

# Decision-Making with Deliberation: Meta-reviewing as a Document-grounded Dialogue

Sukannya Purkayastha<sup>1</sup>, Nils Dycke<sup>1</sup>, Anne Lauscher<sup>2</sup>, Iryna Gurevych<sup>1</sup>

<sup>1</sup> Ubiquitous Knowledge Processing Lab (UKP Lab),

Department of Computer Science, Technical University of Darmstadt and  
National Research Center for Applied Cybersecurity ATHENE, Germany

[www.ukp.tu-darmstadt.de](http://www.ukp.tu-darmstadt.de)

<sup>2</sup> Data Science Group, University of Hamburg

## Abstract

Meta-reviewing is a pivotal stage in the peer-review process, serving as the final step in determining whether a paper is recommended for acceptance. Prior research on meta-reviewing has treated this as a summarization problem over review reports. However, complementary to this perspective, meta-reviewing is a decision-making process that requires weighing reviewer arguments and placing them within a broader context. Prior research has demonstrated that decision-makers can be effectively assisted in such scenarios via dialogue agents. In line with this framing, we explore the practical challenges for realizing dialog agents that can effectively assist meta-reviewers. Concretely, we first address the issue of data scarcity for training dialogue agents by generating synthetic data using Large Language Models (LLMs) based on a self-refinement strategy to improve the relevance of these dialogues to expert domains. Our experiments demonstrate that this method produces higher-quality synthetic data and can serve as a valuable resource towards training meta-reviewing assistants. Subsequently, we utilize this data to train dialogue agents tailored for meta-reviewing and find that these agents outperform *off-the-shelf* LLM-based assistants for this task. Finally, we apply our agents in real-world meta-reviewing scenarios and confirm their effectiveness in enhancing the efficiency of meta-reviewing.<sup>1</sup>

## 1 Introduction

Peer review is the cornerstone of academic quality control across all scientific disciplines (Ware and Mabe, 2009). In a typical peer review process, a group of experts (the *reviewers*) assesses the paper at hand, summarizes their opinion in written reports (the *reviews*), and then a meta-reviewer weighs the presented arguments to decide to *accept*

<sup>1</sup>Code available at: <https://github.com/UKPLab/eacl2026-meta-review-as-dialog>

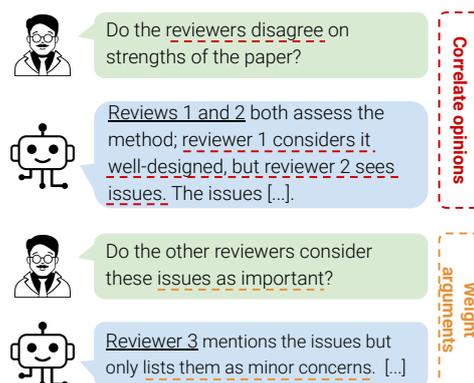


Figure 1: Illustration of the process of *meta-reviewing as a dialogue*. Dialogues include requests to summarize opinions, weight arguments, and contextualize them.

or *reject* the paper for publication resulting in a *meta-review* report. This task requires significant effort and expertise by the meta-reviewer, as they need to match arguments, weigh different opinions, and contextualize the reviews within the broader scientific field. The growing load of research output (Chen et al., 2022; Krenn et al., 2023) puts an additional strain on meta-reviewers requiring to arrive at reliable decisions *quickly*. While the assistance of reviewers through means of natural language processing (NLP) has received considerable attention (e.g., Dycke et al., 2023; Kuznetsov et al., 2024), so far, the support for the decision-making process of meta-reviewers is under-investigated despite the apparent practical needs.

Meta-reviewing has traditionally been framed as a multi-document summarization task (Shen et al., 2022; Li et al., 2023a; Zeng et al., 2025), aimed at generating meta-review reports by condensing the underlying reviews. While these summarization methods can aid meta-reviewers by simplifying the writing process, they often fall short of directly supporting decision-making. A more application-driven and human-centered approach to decision-making remains an important gap.

This paper advocates a new angle on assisting meta-reviewers that focuses on supporting the underlying *decision-making process* to address these needs. For human decision-making and information assimilation over source materials, dialogue assistance has proven effective in many domains (Lever and Sénéchal, 2011; Golinkoff et al., 2019). Following up on a recent proposal by Balke et al. (2023), we explore the practical implications of developing a dialogue agent that assists the meta-reviewer during their step-wise decision-making process by collaboratively reflecting on the underlying peer reviews. Figure 1 illustrates an example of such a dialogue. Ultimately, we investigate the following research questions. (RQ1) **What requirements does the meta-reviewing domain pose on a dialogue agent?** (RQ2) **How can we develop such a dialogue agent to account for these requirements?** (RQ3) **Is a dialogue agent practically useful during meta-reviewing?**

Developing a meta-reviewing decision-making dialogue assistant presents unique challenges. Unlike prior work in other domains on general topics (Wang et al., 2019; Farag et al., 2022), a meta-reviewing dialogue agent must engage with highly technical topics and ensure its responses are strictly grounded in the underlying reviews. As shown in our experiments in Section 3, off-the-shelf Large Language Models (LLMs) are inadequate for this task, as they fail to meet both these requirements. To bridge this gap, dedicated fine-tuning is crucial to tailor the models for technical grounding and decision-making tasks. Additionally, nuanced evaluation methods are required to ensure their reliability. However, fine-tuning and evaluating such agents is costly, as it demands significant effort to generate high-quality human dialogue data.

To address these issues we develop a dedicated fine-tuning and evaluation suite for meta-reviewing dialogue agents. First, we tackle the lack of training data by leveraging LLMs for **synthetic data generation**, which has shown promise across various fields (e.g., Møller et al., 2023; Chen et al., 2023). Specifically, we simulate meta-review decision-making dialogues with LLMs such as ChatGPT (OpenAI, 2026), LLaMa (Touvron et al., 2023), Mistral (Jiang et al., 2023), and Mixtral (Jiang et al., 2024). Through human and automated evaluations, we find that out-of-the-box LLMs generate dialogues that are unspecific and lack diversity. Therefore, we devise a new self-refinement strategy that addresses these challenges in the meta-

reviewing domain but also show its generality for other decision-making domains. Afterwards, in Section 4 we utilize the synthetic dataset to **fine-tune dialogue agents** for meta-reviewing assistance and propose an evaluation strategy. Both automatic and human evaluations show that smaller, fine-tuned open-source models outperform larger, closed zero-shot LLMs like ChatGPT in generating relevant and high-quality responses. Finally, in Section 5, we **deploy the best-performing dialogue agent** within a human user study demonstrating the effectiveness of such a dialogue agent in time-efficient meta-reviewing. Our results show that dialogue agents can effectively **reduce the meta-reviewing time by almost up to 50%**, while enabling the creation of more comprehensive and detailed meta-review reports, compared to when no agent is used. This demonstrates the potential of dialogue agents to significantly enhance both the quality and efficiency of the meta-reviewing process.<sup>2</sup>

In summary, our contributions are three-fold: (1) We present the first study on meta-reviewing as a dialogue to-date and propose extensive measures towards developing dialogue agents for this scenario (RQ1). (2) We address the data scarcity issue for training dialogue agents via synthetic data generation using LLMs. We further propose a self-refinement strategy to enhance the generated dialogues using dialogue quality metrics as feedback (RQ2). (3) We utilize the synthetically generated data for training dialogue agents tailored to the meta-reviewing scenario and evaluate the utility of these agents in real-world meta-reviewing (RQ3).

## 2 Meta-reviewing as a Dialogue

Existing NLP literature treats meta-reviewing exclusively as a summarization problem (Shen et al., 2022; Li et al., 2023a; Zeng et al., 2025). While meta-reviewing involves summarization steps, it is primarily a decision-making process where the human meta-reviewer has to arrive at an overall recommendation on the paper based on the review reports. This decision-making process requires complex reasoning, including weighing, contextualizing, and meta-reasoning on the reviews. For instance, meta-reviewers need to adjust for reviewer uncertainty, biases, and review quality or need to consider the related work and state of the field. Summaries can assist in organizing infor-

<sup>2</sup>Overall workflow of this work is in Appendix A.22

mation but may offer limited direct support for decision-making in meta-reviewing. This is because humans can be influenced by readily presented recommendations, a phenomenon known as *automation bias* (Goddard et al., 2014; Schemmer et al., 2022). In combination with the fact that LLMs sporadically introduce biases and errors in the summaries, this may impact human agency over the process (Donker, 2023; Schintler et al., 2023). Moreover, current ACL policies discourage the use of LLMs as ghostwriters during meta-reviewing (Boyd-Graber et al., 2023).

In response to these issues, we propose a new *human-centered* perspective on meta-reviewing assistance by framing meta-reviewing as a decision-making process that benefits from assistance through a dialogue agent. This concept has shown promise in related domains (Wang et al., 2019; Farag et al., 2022) and in a piloting study by Balke et al. (2023), who simulate this scenario between a human meta-reviewer and a human dialogue assistant. More formally, we define meta-reviewing as a document-grounded decision-making dialogue where a dialogue agent provides information on the underlying reviews, and the meta-reviewer gains insights to move closer to their decision with each turn. In line with the ACL guidelines, such a dialogue agent should *not* pro-actively bring up new facts but act on human inquiries to ensure full human agency. While the agent may still occasionally introduce biases or errors, the meta-reviewer can verify the information and enhance their understanding through iterative refinement (Ke et al., 2024). Specifically, following Balke et al. (2023), the dialogue agent should *ground* their answers in the reviews without hallucinating or obfuscating information. Secondly, the answers should be highly *technically specific*, avoiding generic and unrelated statements. We refer to these two desiderata as GROUNDEDNESS and SPECIFICITY.

Our experiments in Section §3 demonstrate the difficulty of this task. Out-of-the-box LLMs generate answers that fail on both desiderata. Thus, dedicated fine-tuning and nuanced evaluation of these agents are essential. However, data in this domain is costly as human experts and considerable effort is required to generate human-to-human example dialogues as proposed in Balke et al. (2023). We explore these challenges in the remainder of this work.

### 3 Synthetic Data Generation for Fine-tuning Meta-review Dialogue Agents

Developing effective meta-reviewing dialogue agents necessitates dedicated fine-tuning, as off-the-shelf LLMs fall short of this specialized task (cf. Sec §4). This fine-tuning process relies on high-quality training data tailored to meta-reviewing assistance. However, obtaining expert annotators and producing human gold data is resource-intensive since crowd-sourcing is not a viable option (Stasaski et al., 2020). To address these limitations, leveraging LLMs to generate synthetic data has become a widely adopted approach in expert-driven domains such as education (Macina et al., 2023) and biomedicine (Smolyak et al., 2024). Following this trend, we employ LLMs to *synthetically generate* a dataset of meta-reviewing dialogues, building on their proven utility in various tasks (Møller et al., 2023; Chen et al., 2023). However, meta-reviewing dialogues are particularly challenging to synthesize, as they need to account for the previously mentioned desiderata of GROUNDEDNESS and SPECIFICITY. Our initial experiments show that out-of-the-box LLMs fail to generate dialogues that meet these criteria. To overcome this, we implement a self-refinement-based approach for data generation. In the following, we present our data generation approach (§3.1), assess dataset quality (§3.3), fine-tune models (§4), and conduct a within-subject evaluation (§5).

#### 3.1 Approach

We propose ReMuSE (**R**eward-based **M**ulti-aspect **S**elf-**E**dit**E**ng based on the self-refinement paradigm (Madaan et al., 2023; Chen et al., 2024b). In this paradigm, an LLM  $\mathcal{M}$  first produces the initial output  $o$ . Then, feedback  $f$  is obtained on the output via  $\mathcal{M}$  or other feedback mechanisms. The model  $\mathcal{M}$  is then prompted to correct the initial output  $o$  based on  $f$ . All prior approaches (Madaan et al., 2023; Chen et al., 2024b) either improve a single aspect of the output only, or use ground truth in-context exemplars to steer the generations towards a particular format. These factors restrict the diversity of the corrections to a small set of specific changes rather than profound revisions (Wang et al., 2024). However, as our experiments show (cf. Section §3.3), meta-reviewing dialogue generation gains from incorporating multiple aspects during the feedback loop of self-refinement. In this

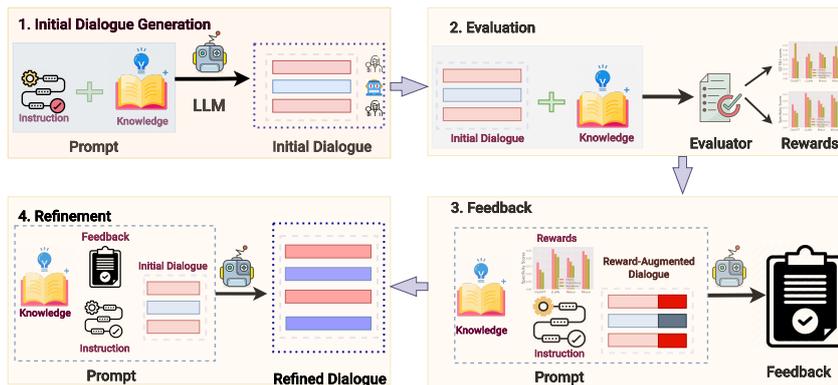


Figure 2: Overview of our **Reward-based Multi-aspect Self Editing (ReMuSE)** method. ReMuSE consists of four steps: 1. **Initial Dialogue Generation** in which we prompt an LLM with relevant documents (paper reviews) and instructions, 2. **Evaluation** of the dialogues by computing one or multiple measures (rewards), 3. Natural language **Feedback** Generation based on the computed rewards, 4. **Self-Refinement** of the dialogues based on the feedback.

work, we thus propose a zero-shot refinement strategy that makes use of scoring functions to calculate dialogue quality metrics. These *rewards* are used as feedback to improve the initial output across multiple dimensions.

ReMuSE generates dialogues based on a knowledge source, e.g. the collection of reviews. Given a knowledge source  $\mathcal{K}$  and an LLM  $\mathcal{M}$  ReMuSE operates in four steps (cf. Figure 2):

**I. Initial Dialogue Generation** We prompt  $\mathcal{M}$  with knowledge source  $\mathcal{K}$  to generate an initial dialogue  $\mathcal{D}$  in a zero-shot fashion.

**II. Evaluation** The generated dialogue is then evaluated to obtain rewards using an **evaluator**. The evaluator consists of various scoring functions. For example, one of the functions in our experiment calculates the token overlap-based groundedness score of utterance,  $u$  with respect to  $\mathcal{K}$ . These rewards are concatenated to every utterance of  $\mathcal{D}$ , producing the reward-augmented dialogue  $\mathcal{D}'$ .<sup>3</sup>

**III. Feedback**  $\mathcal{D}'$  and  $\mathcal{K}$  are fed into  $\mathcal{M}$  to obtain natural language feedback  $\mathcal{F}$  on improving the rewards. Following Madaan et al. (2023), we use the phrase ‘actionable feedback’ in the prompt to obtain a concrete plan to improve the scores.

**IV. Refinement** Finally, we feed the knowledge  $\mathcal{K}$ , the feedback  $\mathcal{F}$ , and the initial dialogue  $\mathcal{D}$  to  $\mathcal{M}$ . The model is prompted to refine the dialogues based on the feedback provided.

### 3.2 Dataset Generation

We generate a dataset of meta-reviewing dialogues for fine-tuning dialogue agents. For this purpose,

<sup>3</sup>The evaluator is completely independent of  $\mathcal{M}$ .

we test several models and perform automated and human evaluations to select the best configuration.

**Models** We use the following models namely, GPT 4o (dubbed ‘ChatGPT’), LLaMa-2 13B (Touvron et al., 2023), Mistral 7B (Jiang et al., 2023), and Mixtral (Jiang et al., 2024).<sup>4</sup>

**Underlying Data** To generate a dialogue, ReMuSE receives a paper and its reviews as the input. As a basis to generate a dataset, we rely on the ORSUM (Zeng et al., 2025) dataset consisting of reviews and meta-reviews for papers from 39 conferences collected using OpenReview.<sup>5</sup> We sample papers with exactly three reviews to fit most LLMs’ context windows resulting in 3064 dialogues.

**Reward Metrics** In the absence of ground truth dialogues to assess the quality of the synthetic dialogues, we use three reference-free metrics as rewards to assess the SPECIFICITY and GROUNDEDNESS of generated dialogues. We use three feedback metrics as rewards to assess the specificity and groundedness of generated dialogues. **Knowledge Precision (K-Prec)** measures token overlap between an utterance and its knowledge source (Adlakha et al., 2024). Because K-Precision cannot capture the semantic similarity between utterances, we additionally employ a semantics-based metric, **Q2-score** (Honovich et al., 2021). This metric evaluates faithfulness via question generation, answer generation, and natural language inference, leveraging pre-trained T5 (Raffel et al., 2020), ALBERT-

<sup>4</sup>Details about the models and prompts in §A.1 and §A.3.

<sup>5</sup>There are no explicit licenses attached to OpenReview for third-party data processing. We use the dataset from the author’s publicly available codebase.

Large (Lan et al., 2020), and DeBERTa-Large-MNLI (He et al., 2021) models, respectively, to capture dialogue groundedness. **Specificity** scores the level of detail in utterances using Specifiteller (Li and Nenkova, 2015).

**Resulting Datasets** Using each LLM, we generate synthetic dialogues. Due to the associated API costs for using closed LLMs, we follow previous works (Bian et al., 2024; Chen et al., 2024a) that rely on a lower number of examples (e.g., 100) for evaluation. Consequently, we sample 200 examples to analyze the performance of ReMuSE and publicly release **3064** meta-reviewing dialogues generated with the best ReMuSE configuration.

**Baselines.** We consider several baselines to compare our method against. (i) **Zero-shot:** LLM is prompted in a 1-pass setup to generate dialogues. (ii) **Self-Refine** (Madaan et al., 2023): the LLM is prompted to refine the generations for the same number of iterations as in our approach, ReMuSE. (iii) **Self-Consistency and Refinement** (Wang et al., 2023; Chen et al., 2024a): The LLM first refines multiple sampled candidates and then the best candidate is selected by a majority vote on the reward criteria.

### 3.3 Dataset Evaluation

We discuss various dialogue quality metrics employed in automated and human evaluation reflecting the desiderata for GROUNDEDNESS and SPECIFICITY. We compare the synthetically generated dialogues to a human-generated dialogue dataset that we use as evaluation data for comparison along the various evaluation setups.

#### 3.3.1 Human Dialogue Dataset

We compare our synthetically generated dialogues with real human meta-reviewing decision-making dialogues using the *Dagstuhl* dataset (Balke et al., 2023). In this dataset, senior and junior scientists acted as meta-reviewers and assistants, respectively, in a wizard-of-oz setup, where the meta-reviewers used dialogue agents to gather information and decide on paper acceptance. The dataset is private and consists of 16 dialogues corresponding to 4 research papers.<sup>6</sup>

#### 3.3.2 Automatic Evaluation

**Setup** We compute K-Prec, Q2-score, Specificity automatically given the knowledge source  $\mathcal{K}$ , and

Model	Method	K-Prec	Q2-F1	Q2-NL	Spec.
ChatGPT	Zero-Shot	0.34 <sub>.01</sub>	0.17 <sub>.01</sub>	0.17 <sub>.01</sub>	0.46 <sub>.13</sub>
LLaMa	Zero-Shot	0.37 <sub>.01</sub>	0.15 <sub>.13</sub>	0.16 <sub>.01</sub>	0.46 <sub>.01</sub>
Mistral	Zero-Shot	0.38 <sub>.01</sub>	0.15 <sub>.02</sub>	0.16 <sub>.01</sub>	0.41 <sub>.02</sub>
Mixtral	Zero-Shot	0.38 <sub>.01</sub>	0.17 <sub>.01</sub>	0.19 <sub>.02</sub>	0.42 <sub>.02</sub>
ChatGPT	Self Ref.	0.46 <sub>.02</sub>	0.19 <sub>.01</sub>	0.19 <sub>.01</sub>	0.52 <sub>.02</sub>
LLaMa	Self Ref.	0.48 <sub>.02</sub>	0.18 <sub>.02</sub>	0.19 <sub>.01</sub>	0.53 <sub>.02</sub>
Mistral	Self Ref.	0.49 <sub>.01</sub>	0.19 <sub>.01</sub>	0.20 <sub>.01</sub>	0.50 <sub>.02</sub>
Mixtral	Self Ref.	0.50 <sub>.01</sub>	0.20 <sub>.01</sub>	0.21 <sub>.02</sub>	0.51 <sub>.02</sub>
ChatGPT	Self Con. + Ref.	0.52 <sub>.02</sub>	0.21 <sub>.02</sub>	0.22 <sub>.02</sub>	0.55 <sub>.02</sub>
LLaMa	Self Con. + Ref.	0.53 <sub>.02</sub>	0.21 <sub>.01</sub>	0.22 <sub>.02</sub>	0.56 <sub>.02</sub>
Mistral	Self Con. + Ref.	0.55 <sub>.01</sub>	0.23 <sub>.02</sub>	0.25 <sub>.01</sub>	0.58 <sub>.02</sub>
Mixtral	Self Con. + Ref.	0.56 <sub>.01</sub>	0.24 <sub>.01</sub>	0.26 <sub>.02</sub>	0.59 <sub>.02</sub>
ChatGPT	ReMuSE	0.76 <sub>.12</sub>	0.24 <sub>.02</sub>	0.24 <sub>.18</sub>	0.61 <sub>.02</sub>
LLaMa	ReMuSE	0.74 <sub>.02</sub>	0.23 <sub>.02</sub>	0.24 <sub>.02</sub>	0.63 <sub>.02</sub>
Mistral	ReMuSE	0.76 <sub>.01</sub>	0.30 <sub>.02</sub>	0.33 <sub>.02</sub>	0.65 <sub>.22</sub>
Mixtral	ReMuSE	<b>0.78<sub>.01</sub></b>	<b>0.32<sub>.01</sub></b>	<b>0.34<sub>.02</sub></b>	<b>0.72<sub>.02</sub></b>

Table 1: Performance with and without the refinement methods (Self-Consistency, Self-Refine, ReMuSE) in terms of K-Prec, Q2-F1, Q2-NLI, and Specificity (Spec) averaged over three prompts with standard deviation.

a dialogue utterance  $u$ , part of the generated dialog. These metrics are also employed as **rewards** within ReMuSE. The evaluation is performed on an utterance level. The intuition behind using the same measures for automated evaluation and for refinement is to understand whether ReMuSE can, in fact, guide the LLMs towards better generations reflected in exactly these dimensions we optimize for. We complement the automatic evaluation with our human evaluation to verify the improved quality.

For the final evaluation, we compute the K-Prec and Q2-scores only over the dialogue agent’s utterances since we require only the dialogue agent to be faithful. The meta-reviewer (as an information seeker) need not fulfill such criteria since they can bring up new facts independent of the knowledge source at any point in time (Dziri et al., 2022). Nevertheless, all the utterances are supposed to be highly specific. This stems from our observation that 65.24% utterances in the human dataset (cf. §3.2) have a score of 0.5 or above, indicating a significant presence of technical details. Hence, we measure the Specificity score over all utterances.

**Overall results and comparison to human dialogues** We compare LLM performance across multiple baselines in Table 1. **ReMuSE consistently improves all models and metrics**, with Mixtral achieving the best results. The largest gains are in **K-Precision** (+105%), with substantial improvements in **Q2-F1** (~88%), **Q2-NLI** (~79%), and **specificity** (~72%). **Self-Refinement** offers modest improvements, reflecting LLM limitations in self-correction (Huang et al., 2024). Self-Consistency with refinement improves performance of the models, however the majority

<sup>6</sup>Additional details about the dataset is in §A.2

Method	Coop.	Coh.	Eng.	Plau.
Zero-Shot	2.35	2.23	2.26	2.45
Self Ref.	2.70	2.45	2.50	2.75
Self Con. + Ref.	3.00	2.60	2.70	3.05
ReMuSE	<b>3.37</b>	2.79	2.85	<b>3.58</b>

Table 2: Human evaluation of the best performing model, Mixtral across different refinement methods. Scores are on a scale of 1 (worst) to 4 (best).

vote mainly selects the longest dialogue reflecting length bias in LLMs (Hu et al., 2025).<sup>78</sup> We further assess **dialogue modeling** with the best performing model, Mixtral, dividing meta-review dialogues into **evidence gathering, conflict resolution, and decision-making stages** following prior approaches on task-specific deliberation (Priya et al., 2025). ReMuSE significantly outperforms other strategies across stages, e.g.,  $\sim 77\%$  vs.  $28\%$  for conflict resolution, demonstrating its effectiveness in structured meta-review dialogues.<sup>910</sup>

In order to compare **human vs synthetic dialogues**, we analyze K-Prec, Q2, and Specificity scores for all utterances in human and synthetic meta-reviewing dialogue datasets, along with {2,3,4}-gram token diversities ( $\{n\}$ -gram vocabulary size) for meta-reviewer utterances in Fig 4. LLM-based dialogue agents outperform humans in Q2 and K-Prec scores (cf. Figures 4c and 4d), demonstrating superior groundedness to the knowledge source. Human meta-reviewers exhibit lower groundedness, possibly due to off-topic divergence: for instance, in one of the dialogues, the meta-reviewer discusses datasets available for a related task not mentioned in the reviews. Regarding Specificity (Fig 4b), LLM-based agents perform comparable to humans, indicating their feasibility as meta-reviewing assistants. However, a significant gap exists between real and LLM meta-reviewers in terms of token diversity and specificity (cf. Fig 4a and 4b), suggesting LLMs still lack the expert knowledge for simulating real-world meta-reviewers.

### 3.3.3 Human Evaluation

**Setup** Since using identical metrics for rewards and automated evaluation can be questioned for coincidental correlation, we further run a human evaluation. Following (Dziri et al., 2022; Wu

<sup>7</sup>Table 7 reports results for ReMuSE with all rewards.

<sup>8</sup>§A.5 for ablations and §A.4 for cross-domain analysis.

<sup>9</sup>Full experimental details in Sec §A.16. Sec §A.23 lists dialogue excerpts exemplifying these stages.

<sup>10</sup>Further results using other open-source models in Sec §A.20

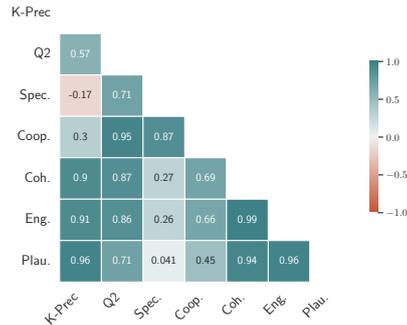


Figure 3: Correlation between human and automated evaluation metrics.

et al., 2022b), we let humans evaluate the generated dialogs according to the following criteria: **Co-operativeness** evaluates the helpfulness of the dialogue agent’s response to the query, **Coherence** specifies consistency of the overall dialogue, and **Engagingness** implies how engagingly the dialogue agent takes the discussion forward. Moreover, akin to faithfulness discussed before, we introduce a **Plausibility** score to measure the groundedness of the dialogue agent’s responses to the reviews. All the dialogues are rated on a Likert scale of 1 – 4 and the ratings are averaged over each metric as done in prior works (Rashkin et al., 2021).<sup>11</sup> **Overall results and comparison to human dialogues** Four Ph.D. students with NLP backgrounds evaluated 70 dialogues, split between the different baseline models and ReMuSE (post-refinement) with the best performing model, Mixtral using human evaluation metrics on a 1-4 Likert scale. Krippendorff’s  $\alpha$  for the dialogues showed substantial agreement: co-operativeness (0.62), coherence (0.74), engagingness (0.74), and plausibility (0.76). We report the results in Table 2. We observe that dialogues generated with ReMuSE are adjudged to be the best across all metrics. We observe a similar ranking of baselines as in our automated evaluation with zero-shot being the least effective method for dialogue generation. We further conducted correlation analysis (cf. Fig 3, §A.13) of the automated and human evaluation metrics which revealed strong alignment between the metrics with similar objective (e.g., K-Prec and Plausibility), validating our evaluation approach.

Three NLP Ph.D. students evaluated the **synthetic and human-generated dialogues** on a 1–4 scale (Sec. §3.3.3). Krippendorff’s  $\alpha$  for co-operativeness, coherence, engagingness, and plau-

<sup>11</sup>The instructions to the annotators are in §A.10

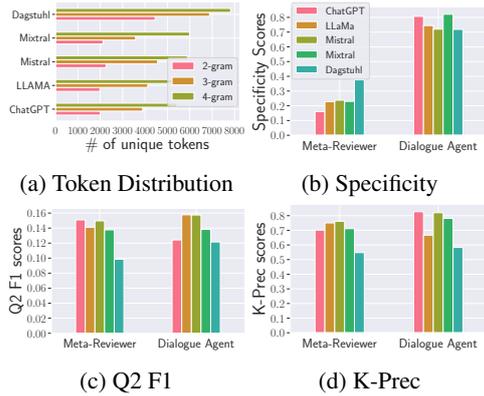


Figure 4: Comparison of utterances in human and synthetically generated dialogues in terms of (a) Token Distribution, (b) Specificity, (c) Q2 F1, and (d) K-Prec.

Setup	Coop.	Coh.	Eng.	Plau.
Dagstuhl	3.50	3.10	<b>3.10</b>	2.95
ReMuSE	<b>3.52**</b>	<b>3.86**</b>	3.00**	<b>3.10**</b>

Table 3: Human evaluation of dialogues from the human dialogue dataset (*Dagstuhl*) and our proposed ReMuSE. The evaluation scale ranges from 1 (worst) - 4 (best). \*\*( $p < .001$ ) statistically significant results

sibility were 0.68, 0.69, 0.70, and 0.74, respectively. As shown in Table 3, synthetic dialogues scored higher for co-operativeness, coherence, and plausibility, while human dialogues scored higher for engagingness. Lower human ratings may stem from off-topic responses and grammatical errors, whereas engagingness benefited from diverse meta-reviewer questions producing longer, detailed responses, consistent with prior findings (e.g., Wu et al., 2022b).

We evaluate generated dialogues using automated metrics for GROUNDEDNESS and SPECIFICITY, and human assessments for flow and naturalness. Our results show LLMs are promising as meta-reviewing assistants, with high faithfulness and specificity, though low token diversity highlights room for improvement as meta-reviewers. Section §4 discusses how fine-tuning can address these limitations.

## 4 Dialogue Agents for Meta-Reviewing Assistance

With the synthetically generated dataset for the best-performing ReMuSE configuration, we now shift our focus to the task of finetuning dialogue agents for meta-reviewing assistance.<sup>12</sup>

<sup>12</sup>See §A.6 and §A.24 for the synthetic dataset and example dialogues.

Models	$(y, \hat{y})$		$(y, \mathcal{K})$		
	BL.	BT Sc.	K-Prec	Q2-F1	Q2-NL
OPT	15.4	54.5	57.6	36.2	40.2
T5	18.4	58.1	63.4	39.1	45.4
Flan-T5	<b>20.4</b>	<b>60.3</b>	<b>68.2</b>	<b>42.2</b>	<b>47.3</b>
ChatGPT	12.1	45.0	40.1	23.1	22.4
ReMuSE	14.5	49.0	67.6	42.2	47.1

Table 4: Performance of models in the response generation task. BLEU (BL.) and BERT Score (BT Sc.) measure the similarity of the response with the ground truth. K-Prec and Q2 scores evaluate the factual consistency of the response with the Knowledge Source.

Models	Fl.	Rel.	Faith.	Obj.	Help.
OPT	3.82	3.84	4.04	4.06	4.21
T5	3.91	3.98	4.12	4.21	4.32
Flan-T5	<b>4.21</b>	<b>4.18</b>	<b>4.58</b>	<b>4.62</b>	<b>4.71</b>
ChatGPT	3.80	3.80	3.24	3.87	3.82
ReMuSE	3.92	3.94	4.11	4.11	4.24

Table 5: Human evaluation of model responses on Fluency, Relevance, Faithfulness, Objectivity, and Helpfulness scored 1 (worst) to 5 (best).

### 4.1 Experimental Setup

**Task** Given a knowledge source  $\mathcal{K}$  and a dialogue history  $[u_1, u_2, \dots, u_t]$ , the task is to generate a response  $u_{t+1}$  grounded in  $\mathcal{K}$ . The title and reviews for every paper serve as the knowledge source.

**Models** Following Daheim et al. (2024), we select Flan-T5 (Longpre et al., 2023) (x1), T5 (Raffel et al., 2020) (3B), and OPT (Zhang et al., 2022) (2.7B) for experiments. For further comparison, we also prompt ChatGPT in a 0-shot setup and use the best model from the ReMuSE setup, Mistral.<sup>13</sup>

**Data and Evaluation** We split the 3064 meta-reviewing dialogues having an average of 12.24 utterances per dialogue in our dataset into 60% for training, 20% for validation, and 20% for testing, respectively. In terms of evaluation, following Daheim et al. (2024), we use automated metrics - SacreBLEU (Post, 2018) and BERTScore (Zhang et al., 2020) to measure response similarity to ground truth. We also report K-Prec and Q2 scores to assess faithfulness to  $\mathcal{K}$ . Additionally, following Rashkin et al. (2021), we perform human evaluation on fluency, relevance, faithfulness, and objectivity, plus a helpfulness score reflecting decision-making utility. All human ratings use a 1–5 scale, and the ratings are averaged over each metric.

### 4.2 Results and Discussion

**Automated Evaluation** Table 4 lists model performances. Flan-T5 achieves the highest

<sup>13</sup>The hyper-parameter details are provided in §A.7

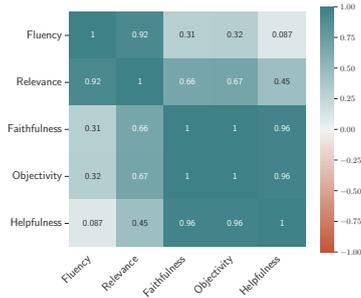


Figure 5: Correlation of the human evaluation metrics. We observe the strongest correlation between Helpfulness, Faithfulness, and Objectivity.

K-Prec (68.2), consistent with prior work (Daheim et al., 2024), while ChatGPT performs significantly lower (K-Prec 42.1), highlighting the importance of high-quality task-specific supervision. We show that ReMuSE maintains strong faithfulness with a K-Prec of 67.6, though its BLEU and BERT scores are lower than fine-tuned models because it is optimized for faithful and diverse generation rather than surface-level alignment. In contrast, fine-tuned models benefit from task-specific supervision, improving alignment with reference responses (Hsieh et al., 2023; Fu et al., 2024). We further show in §A.17 that even with only 10% of the data, models trained on ReMuSE significantly outperform zero-shot variants, demonstrating the high quality of the generated dialogues.

**Human Evaluation** Following Hossain et al. (2025), we evaluate dialogue agents on their ability to assist meta-reviewers by rating responses to key queries on contributions, strengths, weaknesses, improvements, and literature review. Two senior NLP Ph.D. students with over three years of reviewing experience rated 60 responses per model on a 1–5 Likert scale (Purkayastha et al., 2023).<sup>14</sup> Krippendorff’s  $\alpha$  for fluency, relevance, faithfulness, objectivity, and helpfulness are 0.58, 0.64, 0.69, 0.66, and 0.74, respectively, substantial given task subjectivity (Kennard et al., 2022).<sup>15</sup> Table 5 shows that Flan-T5 scores highest, while ChatGPT scores lowest on faithfulness and objectivity. ReMuSE outperforms ChatGPT but lags behind fine-tuned models, likely due to its diversity-driven objectives and verbose responses, often ending with follow-up questions. Helpfulness correlates strongly with faithfulness and objectivity (Fig. 5), supporting the

<sup>14</sup>See §A.11 for annotator instructions.

<sup>15</sup>Annotator disagreements are analyzed in §A.14

Metrics	w/o Dial. Agent	w. Dial. Agent	Auto Meta-Rev
<b>Expert</b>			
Content Relevance	3.42	<b>3.87</b>	3.55
Decision Correctness	<b>0.70</b>	0.62	0.65
Coverage	3.21	<b>3.76</b>	3.40
<b>Non-expert</b>			
Content Relevance	3.10	<b>3.65</b>	3.32
Decision Correctness	0.58	<b>0.60</b>	0.59
Coverage	2.95	<b>3.50</b>	3.20

Table 6: Human evaluation of meta-reviews written by experts and non-experts, with and without dialogue agent assistance, and those generated automatically (Auto Meta-Rev). Content Relevance and Coverage are rated on a scale of 1–4. Decision Correctness is binary (0/1).

quality of our synthetic training data.<sup>16</sup>

## 5 Deploying Meta-Reviewing Dialogue Agents: A Within-subject Study

We analyze the trained dialogue agent’s effectiveness in real-world meta-reviewing, focusing on time and meta-review quality.

**Setup** We conduct a within-subject controlled experiment where participants write meta-reviews with and without dialogue agent support.<sup>17</sup> Following prior work on user studies in peer review (Hossain et al., 2025), we ask fifteen Ph.D. students with peer review experience to evaluate three papers each under both conditions (45 papers total).<sup>18</sup> Following Shen et al. (2022), we assess meta-review quality using (i) *Content Relevance* (alignment with review aspects, rated 1-4), (ii) *Decision Correctness* (alignment with the gold decision, binary 0/1), and (iii) *Coverage* (review diversity, rated 1-4). Two senior Ph.D. students evaluate the meta-reviews alongside an automatically generated meta-review across the chosen metrics. We employ the Level-1 prompting framework introduced in Hossain et al. (2025), wherein the LLM is tasked with generating a meta-review conditioned on the reviews of a specific paper.<sup>19</sup>

**Results.** As shown in Table 6, reviews generated with the dialogue agent scored higher in Content Relevance (3.87 vs. 3.42) and Coverage (3.76 vs. 3.21), with Decision Correctness comparable (0.62 vs. 0.70). The automated meta-review baseline performed better than the no-agent condition (Content Relevance 3.55, Coverage 3.40, Decision

<sup>16</sup>Meta-reviewer question distribution and error analysis of the dialogue agent in Sec §A.8 and Sec §A.9

<sup>17</sup>The interaction interface is in Fig 10, §A.15

<sup>18</sup>The participant details is in §A.12

<sup>19</sup>Operation guidelines for deployment in Sec §A.21.

Correctness 0.65) but below the dialogue-assisted setup. Interaction with the agent took an average of **20 minutes**, compared to **35 minutes** without it, demonstrating improved efficiency without compromising quality. The latency of the agent is 14.4 seconds for 12 turns ( $\sim 1.2$  seconds per turn)

**Case Study with Non-Experts** While meta-reviewing is usually done by experienced researchers, we test whether individuals from interdisciplinary backgrounds can use the tool effectively. Ten PhD students from non-CS fields, with no prior peer-review experience, wrote meta-reviews for three papers each under both setups.<sup>20</sup> Agent-assisted reviews outperformed the no-agent condition in Content Relevance (3.65 vs. 3.10), Coverage (3.50 vs. 2.95), and Decision Correctness (0.60 vs. 0.58). The automated baseline improved over no-agent (3.32, 3.20, 0.59) but was still weaker than agent-assisted reviews. Participants spent an average of **35 minutes** with the agent versus **80 minutes** without it, showing the agent’s efficiency. Dialogues were longest for borderline papers (**12.2 turns**) with the highest conflict resolution coverage (**77%**), demonstrating the agent’s support for structured deliberation and meta-review decision-making (Sec. §A.18).

Paired t-tests and Wilcoxon signed-rank tests with Cohen’s d show robust improvements (Cleophas and Zwinderman, 2010). For 15 PhD students, the dialogue agent increased content relevance by 0.42 and coverage by 0.52, saving 15 minutes. For 10 non-PhD participants, relevance and coverage rose by 0.54 each, with a 45-minute reduction. Gains exceed those of automated meta-reviewing (+0.13–0.22, +0.19–0.25), and are statistically significant, likely generalizing to a larger population.<sup>21</sup>

## 6 Related Work

**Meta-Review Generation** Meta-reviewing is commonly framed as multi-document summarization on the reviews. Wu et al. (2022a) treat reviews, author responses, and reviewer discussions as an argumentative graph to generate comprehensive meta-reviews. Li et al. (2023a) encode entire reviewer conversation threads using a multi-task approach to meta-review generation. Shen et al. (2022) introduce control codes for generating diverse meta-reviews. Zeng et al. (2025) use LLMs in a self-

refinement paradigm based on task-specific checklists. Recently, Hossain et al. (2025) analyzed the performance of LLMs for generating various meta-reviewing aspects. However, we redefine meta-reviewing as a dialogue that aids meta-reviewers in decision-making, rather than automatically generating text, preserving their scientific autonomy.

**AI-assisted decision making** In this work, we frame meta-reviewing as a decision-making process. The role of AI in supporting human decisions has expanded across fields like finance, law, and technology (Aldasoro et al., 2025; Lai et al., 2023), following the collaborative decision-making paradigm (Lin et al., 2024). While Lin et al. (2024) explore reward-based models with fixed solutions, we focus on meta-reviewing, where no single best solution exists, and decisions depend on the meta-reviewer’s expertise.

**Synthetic Dialogue Generation** We fine-tune meta-review agents on synthetically generated, review-grounded dialogues that reflect reasoning behind meta-review decisions. Prior synthetic data approaches (Dai et al., 2022; Bao et al., 2023) are limited by scale. We leverage LLMs as cost-effective curators to produce full-length, knowledge-grounded dialogues with high faithfulness and specificity. Using a zero-shot self-refinement strategy (Madaan et al., 2023), we iteratively enhance multiple aspects of dialogue quality, outperforming prior refinement methods (Wang et al., 2023) and enabling high-quality training data for complex tasks.

## 7 Conclusion

In this work, we explored the potential of dialogue agents to assist meta-reviewers in the decision-making process, offering an alternative to traditional summarization approaches. We addressed the data scarcity challenge with a self-refinement method for generating diverse, faithful decision-making dialogues using LLMs. We found that models fine-tuned on our data significantly outperformed their zero-shot counterparts. Using a controlled experiment, we demonstrated that dialogue agents reduced the meta-reviewing time by almost half while maintaining reviewer autonomy. Our findings highlight the potential of AI-driven tools to improve peer-reviewing efficiency and encourage further research in this area.

<sup>20</sup>Participant details in §A.12

<sup>21</sup>Detailed statistics and effect sizes in §A.19

## Limitations

**Bias** We employ both closed and open-source large language models that are standard for dialogue generation. These models have been criticized for producing harmful or biased content which may reflect in the generated synthetic data. We do not address biases such as socio-demographic factors, given the assumed neutrality of the decision-making domain. Additionally, we do not consider potential bias in the knowledge sources, such as paper reviews in ORSUM, which may include subjective content. Our focus is on generating faithful and diverse dialogues; analyzing the effect of bias in the underlying data is left to future work. We also emphasize that the agent cannot compensate for poor-quality reviews – addressing review quality remains a critical, orthogonal research direction.

**Practical Application** LLMs often lack interpretability, making it difficult to trace how recommendations are generated. We do not aim to improve interpretability in this work, and the responsibility for validating suggestions lies with the meta-reviewer. Accountability is especially important when agent outputs influence decisions. While our agent is designed not to make final judgments, it may occasionally offer flawed suggestions. To mitigate this, reviewer training and clear usage guidelines are essential.

**Scope** Our study focuses on English-language reviews in AI conferences. The limited availability of high-quality, annotated datasets outside well-studied domains—particularly in fields such as the humanities—currently restricts the generalizability and extension of this work to those areas. This data scarcity poses challenges for both model training and rigorous evaluation, thereby constraining cross-domain applicability. Future work could address this limitation by curating or leveraging domain-specific datasets and extending the proposed approach to multilingual settings and underexplored domains, including various subfields within the humanities, to assess its robustness and broader applicability.

**Evaluation** We further note that improvements in “decision correctness” should be interpreted as enhanced *alignment* with historical meta-review outcomes rather than objective accuracy. Real conference decisions are shaped by multiple contextual factors –acceptance quotas, program-chair policies,

and field-specific norms – that are beyond the scope of our work. Our evaluation isolates the deliberative quality of dialogue support rather than predicting actual acceptance decisions. Future work could incorporate contextual priors (e.g., acceptance-rate distributions) to model decision processes more realistically.

## Ethics Statement

In this work, we provide a detailed study of the dialogue generation capabilities of Large language models for meta-reviewing. Consequently, this work has an ethical dimension since it can be perceived as an attempt to replace human meta-reviewers with LLMs. We would like to emphasize that we solely focus on assisting the meta-reviewers in decision making using dialogue agents rather than providing them with any subjective opinions. Unlike previous work, we do not attempt in any way to generate meta-reviews and strongly believe that such a task is a ‘**Human-In-The-Loop**’ process which requires highly skilled researchers (meta-reviewers) operating as humans to make the loop effective.

## Acknowledgements

This work has been funded by the German Research Foundation (DFG) as part of the Research Training Group KRITIS No. GRK 2222, along with the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts, within their joint support of the National Research Center for Applied Cybersecurity ATHENE. We gratefully acknowledge the support of Microsoft with a grant for access to OpenAI GPT models via the Azure cloud (Accelerate Foundation Model Academic Research). The work of Anne Lauscher is funded under the Excellence Strategy of the German Federal Government and the Federal States. The authors acknowledge the support of Schloss Dagstuhl – Leibniz Center for Informatics through the Dagstuhl Seminar ‘22432: *Towards a Unified Model of Scholarly Argumentation*’.

We thank Jonathan Tonglet, Federico Tiblias, Anmol Goel, Indraneil Paul, Aishik Mandal, Luke Bates, Aniket Pramanick, and Chen Cecilia Liu for their help with annotation; Hiba Arnout and Luke Bates for initial feedback on a draft of this paper; and Ming Tsun Lai, Andrew Eliseev, Jens Wala, Letizia Curreri, and Yannick Burmeister for fruitful

discussions and evaluation of the dialogue agent.

## References

- Vaibhav Adlakha, Parishad BehnamGhader, Xing Han Lu, Nicholas Meade, and Siva Reddy. 2024. [Evaluating correctness and faithfulness of instruction-following models for question answering](#). *Transactions of the Association for Computational Linguistics*, 12:681–699.
- I. Aldasoro, L. Gambacorta, A. Korinek, V. Shreeti, and M. Stein. 2025. [Intelligent financial system: How ai is transforming finance](#). *Journal of Financial Stability*, 81:101472. <https://doi.org/10.1016/j.jfs.2025.101472>.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *ArXiv preprint arXiv:2204.05862*.
- Wolf-Tilo Balke, Andreas Vlachos, Davide Ceolin, Midad Alshomary, Nils Dyrce, Sukannya Purkayastha, Iryna Gurevych, Anne Lauscher, and Tilman Beck. 2023. [5.4 scholarly argumentation as a community dialogue](#). *Dagstuhl Reports*, 12(10):202. <https://doi.org/10.4230/DagRep.12.10.175>.
- Jianzhu Bao, Rui Wang, Yasheng Wang, Aixin Sun, Yitong Li, Fei Mi, and Ruifeng Xu. 2023. [A synthetic data generation framework for grounded dialogues](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10866–10882, Toronto, Canada. Association for Computational Linguistics.
- Ning Bian, Xianpei Han, Le Sun, Hongyu Lin, Yaojie Lu, Ben He, Shanshan Jiang, and Bin Dong. 2024. [ChatGPT is a knowledgeable but inexperienced solver: An investigation of commonsense problem in large language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3098–3110, Torino, Italia. ELRA and ICCL.
- Jordan Boyd-Graber, Naoaki Okazaki, and Anna Rogers. 2023. [ACL’23 Peer Review Policies](#). <https://2023.aclweb.org/blog/review-acl23/>. Accessed: 2025-12-02.
- Justin Chen, Swarnadeep Saha, and Mohit Bansal. 2024a. [ReConcile: Round-Table Conference Improves Reasoning via Consensus among Diverse LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7066–7085, Bangkok, Thailand. Association for Computational Linguistics.
- Maximillian Chen, Alexandros Papangelis, Chenyang Tao, Seokhwan Kim, Andy Rosenbaum, Yang Liu, Zhou Yu, and Dilek Hakkani-Tur. 2023. [PLACES: Prompting language models for social conversation synthesis](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 844–868, Dubrovnik, Croatia. Association for Computational Linguistics.
- Pinzhen Chen, Zhicheng Guo, Barry Haddow, and Kenneth Heafield. 2024b. [Iterative Translation Refinement with Large Language Models](#). In *Proceedings of the 25th Annual Conference of the European Association for Machine Translation (Volume 1)*, pages 181–190, Sheffield, UK. European Association for Machine Translation (EAMT).
- Xieling Chen, Haoran Xie, and Xiaohui Tao. 2022. [Vision, status, and research topics of natural language processing](#). *Natural Language Processing Journal*, 1:100001. <https://doi.org/10.1016/j.nlp.2022.100001>.
- Ton J. Cleophas and Aeilko H. Zwinderman. 2010. [Paired Continuous Data \(Paired-t, Wilcoxon\) \(10 Patients\)](#), pages 7–9. Springer Netherlands, Dordrecht. [https://doi.org/10.1007/978-90-481-9519-0\\_3](https://doi.org/10.1007/978-90-481-9519-0_3).
- Nico Daheim, Nouha Dziri, Mrinmaya Sachan, Iryna Gurevych, and Edoardo Ponti. 2024. [Elastic weight removal for faithful and abstractive dialogue generation](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7096–7112, Mexico City, Mexico. Association for Computational Linguistics.
- Zhuyun Dai, Arun Tejasvi Chaganty, Vincent Y. Zhao, Aida Amini, Qazi Mamunur Rashid, Mike Green, and Kelvin Guu. 2022. [Dialog inpainting: Turning documents into dialogs](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 4558–4586. PMLR.
- Tjibbe Donker. 2023. [The dangers of using large language models for peer review](#). *The Lancet Infectious Diseases*, 23(7):781. [https://doi.org/10.1016/S1473-3099\(23\)00290-6](https://doi.org/10.1016/S1473-3099(23)00290-6).
- Nils Dyrce, Ilia Kuznetsov, and Iryna Gurevych. 2023. [NLPeer: A unified resource for the computational study of peer review](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5049–5073, Toronto, Canada. Association for Computational Linguistics.

- Nouha Dziri, Ehsan Kamaloo, Sivan Milton, Omar Zaiane, Mo Yu, Edoardo M. Ponti, and Siva Reddy. 2022. **FaithDial: A Faithful Benchmark for Information-Seeking Dialogue**. *Transactions of the Association for Computational Linguistics*, 10:1473–1490.
- Younna Farag, Charlotte Brand, Jacopo Amidei, Paul Piwek, Tom Stafford, Svetlana Stoyanchev, and Andreas Vlachos. 2022. **Opening up minds with argumentative dialogues**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4569–4582, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xue-Yong Fu, Md Tahmid Rahman Laskar, Elena Khasanova, Cheng Chen, and Shashi Tn. 2024. **Tiny titans: Can smaller large language models punch above their weight in the real world for meeting summarization?** In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 387–394, Mexico City, Mexico. Association for Computational Linguistics.
- Iftah Gamzu, Hila Gonen, Gilad Kutiel, Ran Levy, and Eugene Agichtein. 2021. **Identifying helpful sentences in product reviews**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 678–691, Online. Association for Computational Linguistics.
- Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt. 2014. **Automation bias: Empirical results assessing influencing factors**. *International Journal of Medical Informatics*, 83(5):368–375. <https://doi.org/10.1016/j.ijmedinf.2014.01.001>.
- Roberta Michnick Golinkoff, Erika Hoff, Meredith L Rowe, Catherine S Tamis-LeMonda, and Kathy Hirsh-Pasek. 2019. Language matters: Denying the existence of the 30-million-word gap has serious consequences. *Child development*, 90(3):985–992. <https://doi.org/10.1111/cdev.13128>.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. **DEBERTA: Decoding-Enhanced BERT with Disentangled Attention**. In *9th International Conference on Learning Representations, ICLR 2021, Virtual, May 03-07, 2021*. Openreview.net.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. **The curious case of neural text degeneration**. In *Proceedings of the 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Or Honovich, Leshem Choshen, Roei Aharoni, Ella Neeman, Idan Szpektor, and Omri Abend. 2021. **q<sup>2</sup>: Evaluating factual consistency in knowledge-grounded dialogues via question generation and question answering**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7856–7870, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Eftekhari Hossain, Sanjeev Kumar Sinha, Naman Bansal, R. Alexander Knipper, Souvika Sarkar, John Salvador, Yash Mahajan, Sri Ram Pavan Kumar Guttikonda, Mousumi Akter, Md. Mahadi Hassan, Matthew Freestone, Matthew C. Williams Jr., Dongji Feng, and Santu Karmaker. 2025. **LLMs as meta-reviewers’ assistants: A case study**. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7763–7803, Albuquerque, New Mexico. Association for Computational Linguistics.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. **Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8003–8017, Toronto, Canada. Association for Computational Linguistics.
- Zhengyu Hu, Linxin Song, Jieyu Zhang, Zheyuan Xiao, Tianfu Wang, Zhengyu Chen, Nicholas Jing Yuan, Jianxun Lian, Kaize Ding, and Hui Xiong. 2025. **Explaining length bias in LLM-based preference evaluations**. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 6763–6794, Suzhou, China. Association for Computational Linguistics.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2024. **Large language models cannot self-correct reasoning yet**. In *12th International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 07-11, 2024*. Openreview.net.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, et al. 2023. **Mistral 7b**. *ArXiv preprint arXiv: 2310.06825*.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, et al. 2024. **Mixtral of Experts**. *ArXiv preprint arXiv: 2401.04088*.
- Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, et al. 2025. **Gemma 3 technical report**. *Arxiv preprint arXiv: 2503.19786*.
- Yuhe Ke, Rui Yang, Sui An Lie, Taylor Xin Yi Lim, Yilin Ning, Irene Li, Hairil Rizal Abdullah, Daniel

- Shu Wei Ting, and Nan Liu. 2024. Mitigating cognitive biases in clinical decision-making through multi-agent conversations using large language models: simulation study. *Journal of Medical Internet Research*, 26:e59439. <https://doi.org/10.2196/59439>.
- Neha Nayak Kennard, Tim O’Gorman, Rajarshi Das, Akshay Sharma, Chhandak Bagchi, Matthew Clinton, Pranay Kumar Yelugam, Hamed Zamani, and Andrew McCallum. 2022. **DISAPERE: A dataset for discourse structure in peer review discussions**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1234–1249, Seattle, United States. Association for Computational Linguistics.
- Mario Krenn, Lorenzo Buffoni, Bruno Coutinho, Sagi Eppel, Jacob Gates Foster, Andrew Gritsevskiy, Harlin Lee, Yichao Lu, João P. Moutinho, Nima Sanjabi, Rishi Sonthalia, Ngoc Mai Tran, Francisco Valente, Yangxinyu Xie, Rose Yu, and Michael Kopp. 2023. **Forecasting the future of artificial intelligence with machine learning-based link prediction in an exponentially growing knowledge network**. *Nature Machine Intelligence*, 5(11). <http://dx.doi.org/10.1038/s42256-023-00735-0>.
- Iliia Kuznetsov, Osama Mohammed Afzal, Koen Dercksen, Nils Dycke, Alexander Goldberg, Tom Hope, Dirk Hovy, et al. 2024. **What can natural language processing do for peer review?** *ArXiv preprint arXiv: 2405.06563*.
- Vivian Lai, Chacha Chen, Alison Smith-Renner, Q. Vera Liao, and Chenhao Tan. 2023. **Towards a science of human-ai decision making: An overview of design space in empirical human-subject studies**. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’23*, page 1369–1385, New York, NY, USA. Association for Computing Machinery.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. **ALBERT: A lite BERT for Self-supervised Learning of Language Representations**. In *Proceedings of the 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Rosemary Lever and Monique Sénéchal. 2011. **Discussing stories: On how a dialogic reading intervention improves kindergartners’ oral narrative construction**. *Journal of experimental child psychology*, 108(1):1–24. <https://doi.org/10.1016/j.jecp.2010.07.002>.
- Junyi Jessy Li and Ani Nenkova. 2015. **Fast and accurate prediction of sentence specificity**. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pages 2281–2287. AAAI Press.
- Miao Li, Eduard Hovy, and Jey Lau. 2023a. **Summarizing multiple documents with conversational structure for meta-review generation**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7089–7112, Singapore. Association for Computational Linguistics.
- Siheng Li, Cheng Yang, Yichun Yin, Xinyu Zhu, Zesen Cheng, Lifeng Shang, Xin Jiang, Qun Liu, and Yujia Yang. 2023b. **AutoConv: Automatically generating information-seeking conversations with large language models**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1751–1762, Toronto, Canada. Association for Computational Linguistics.
- Jessy Lin, Nicholas Tomlin, Jacob Andreas, and Jason Eisner. 2024. **Decision-oriented dialogue for human-AI collaboration**. *Transactions of the Association for Computational Linguistics*, 12:892–911.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **RoBERTa: A Robustly Optimized BERT Pretraining Approach**. *Arxiv preprint arXiv: 1907.11692*.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. **The Flan Collection: Designing Data and Methods for Effective Instruction Tuning**. In *Proceedings of the 40th International Conference on Machine Learning, ICML 2023, Honolulu Hawaii USA, July 23 - 29, 2023*. PMLR.
- Edisa Lozić and Benjamin Štular. 2023. **Fluent but not factual: A comparative Analysis of Chatgpt and other AI chatbots’ Proficiency and Originality in Scientific Writing for Humanities**. *Future Internet*, 15(10). <https://doi.org/10.3390/fi15100336>.
- Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023. **MathDial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5602–5621, Singapore. Association for Computational Linguistics.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. **Self-refine: Iterative Refinement with Self-Feedback**. In *Proceedings of the Thirty-seventh Conference on Neural Information Processing Systems, NeurIPS 2023, New Orleans, USA*. Openreview.net.
- Anders Giovanni Møller, Jacob Aarup Dalsgaard, Arianna Pera, and Luca Maria Aiello. 2023. **Is a Prompt**

- and a few samples all you need? using GPT-4 for data augmentation in low-resource classification tasks. *Arxiv preprint arXiv: 2304.13861*.
- OpenAI. 2026. Chatgpt. <https://chat.openai.com/>. Accessed: 2026-01-09.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Proceedings of the Thirty-sixth Conference on Neural Information Processing Systems, NeurIPS 2022, New Orleans, USA (Hybrid), Nov 28- Dec 09, 2022*. Openreview.net.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Priyanshu Priya, Saurav Dudhate, Desai Vishesh Yasheshbhai, and Asif Ekbal. 2025. [We argue to agree: Towards personality-driven argumentation-based negotiation dialogue systems for tourism](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 25504–25536, Suzhou, China. Association for Computational Linguistics.
- Sukannya Purkayastha, Anne Lauscher, and Iryna Gurevych. 2023. [Exploring jiu-jitsu argumentation for writing peer review rebuttals](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14479–14495, Singapore. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21:140:1–140:67. <https://dl.acm.org/doi/10.5555/3455716.3455856>.
- Hannah Rashkin, David Reitter, Gaurav Singh Tomar, and Dipanjan Das. 2021. [Increasing faithfulness in knowledge-grounded dialogue with controllable features](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 704–718, Online. Association for Computational Linguistics.
- Max Schemmer, Niklas Kühl, Carina Benz, and Gerhard Satzger. 2022. [On the influence of explainable ai on automation bias](#). In *Proceedings of the 30th European Conference on Information Systems (ECIS), Timișoara, RO, June 18-24, 2022*. aisel.aisnet.org.
- Laurie A Schintler, Connie L McNeely, and James Witte. 2023. [A Critical Examination of the ethics of AI-Mediated Peer Review](#). *Arxiv preprint arXiv: 2309.12356*.
- Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You, and Luo Si. 2022. [MRcD: A meta-review dataset for structure-controllable text generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2521–2535, Dublin, Ireland. Association for Computational Linguistics.
- Daniel Smolyak, Margrét V Bjarnadóttir, Kenyon Crowley, and Ritu Agarwal. 2024. [Large language models and synthetic health data: progress and prospects](#). *JAMIA Open*, 7(4):ooae114. <https://doi.org/10.1093/jamiaopen/ooae114>.
- Katherine Stasaski, Kimberly Kao, and Marti A. Hearst. 2020. [CIMA: A large open access dialogue dataset for tutoring](#). In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–64, Seattle, WA, USA → Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv preprint arXiv: 2307.09288*.
- Xuewei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. [Persuasion for good: Towards a personalized persuasive dialogue system for social good](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *Proceedings of the 11th International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 01-05, 2023*. Openreview.net.
- Ziqi Wang, Le Hou, Tianjian Lu, Yuexin Wu, Yunxuan Li, Hongkun Yu, and Heng Ji. 2024. [Enabling language models to implicitly learn self-improvement](#). In *12th International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 07-11, 2024*. OpenReview.Net.
- Mark Ware and Michael Mabe. 2009. [An overview of scientific and scholarly journal publishing](#). *The STM report*, 1082:1083. <https://digitalcommons.unl.edu/cgi/viewcontent.cgi?article=1008&context=scholcom>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen,

- Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Po-Cheng Wu, An-Zi Yen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2022a. [Incorporating peer reviews and rebuttal counter-arguments for meta-review generation](#). In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, CIKM '22*, page 2189–2198, New York, NY, USA. Association for Computing Machinery.
- Qingyang Wu, Song Feng, Derek Chen, Sachindra Joshi, Luis Lastras, and Zhou Yu. 2022b. [DG2: Data augmentation through document grounded dialogue generation](#). In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 204–216, Edinburgh, UK. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, et al. 2025. [Qwen3 technical report](#). *Arxiv preprint arXiv: 2505.09388*.
- Qi Zeng, Mankeerat Sidhu, Ansel Blume, Hou Pong Chan, Lu Wang, and Heng Ji. 2025. [Scientific opinion summarization: Paper meta-review generation dataset, methods, and evaluation](#). In *Artificial Intelligence for Research and Democracy: First International Workshop, AI4Research 2024, and 4th International Workshop, DemocrAI 2024, Held in Conjunction with IJCAI 2024, Jeju, South Korea, August 5, 2024, Proceedings*, page 20–38, Berlin, Heidelberg. Springer-Verlag.
- Justine Zhang, Ravi Kumar, Sujith Ravi, and Cristian Danescu-Niculescu-Mizil. 2016. [Conversational flow in Oxford-style debates](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 136–141, San Diego, California. Association for Computational Linguistics.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, et al. 2022. [OPT: Open Pre-trained Transformer Language Models](#). *ArXiv preprint arXiv: 2205.01068*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *Proceedings of the 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

## A Appendix

### A.1 Model Details and Computational Budget

**GPT 4o** We use the chat-tuned variant (dubbed ‘*ChatGPT*’) of the Instruct-GPT class of models from OpenAI (Ouyang et al., 2022).<sup>22</sup> This model was trained with internet data and finetuned using reinforcement learning from human feedback (Bai et al., 2022). The training and dataset details are not publicly available. We use ‘gpt4o’.

**LLaMa-2** It is the upgraded version of the LLaMa family of models (Touvron et al., 2023) featuring a 40% increase in pre-training models’ data and doubled context length compared to its predecessor. We employ the 13B-chat version in our experiments.

**Mistral** This model is trained on the grouped query and sliding window attention that helps focus on words even outside the context window (Jiang et al., 2023). We use the 7B version.

**Mixtral** It is an open-source LLM based on a mixture of experts (Jiang et al., 2024). Each feed-forward block has access to 8 parameter groups and utilizes two of them to process a token.

**Gemma** Gemma is an open-source LLM designed for instruction-following and reasoning tasks (Kamath et al., 2025). We use the 12B-IT version.

**Qwen-3** Qwen-3 is a multilingual LLM capable of handling text and code across multiple languages (Yang et al., 2025). We use the 8B-IT version.

We use the Azure OpenAI service to prompt ChatGPT. The rest of the models are implemented using huggingface transformers (Wolf et al., 2020). Following previous work on data generation with LLMs (Adlakha et al., 2024), we use a high temperature of 0.95 and employ top-p sampling (Holtzman et al., 2020) of 0.95 to avoid sampling repetitions.

#### A.1.1 Computational Budget

We ran all the experiments on Nvidia A100 80GB GPUs. None of the experiments consumed more than 36 hours.

<sup>22</sup><https://chat.openai.com/>

### A.2 Dagstuhl dataset

For the comparison of our synthetically generated dialogues with real human meta-reviewing decision making dialogues we employ the *Dagstuhl* dataset (Balke et al., 2023) kindly provided by the authors. For this dataset, 4 senior and 4 junior scientists acted as meta-reviewers and assistants, respectively, in a wizard-of-oz scenario. The study was conducted on 4 papers selected from the NLPEER corpus (Dycke et al., 2023). The assistants were provided with the paper and reviews in advance, whereas the meta-reviewers were instructed to skim the reviews within 10 minutes. The meta-reviewers were asked to converse with the dialogue agents to gather information and eventually decide on the acceptance of the paper. The paper acceptance decisions for the dataset are displayed in Figure 6. Across the 4 papers, only 1 paper had a unanimous reject decision, while all the other papers received diverse verdicts.

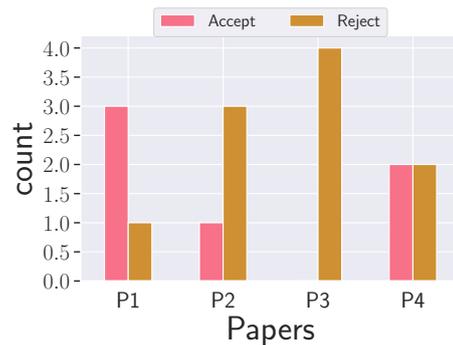


Figure 6: Diversity of Paper Acceptance Decisions in the Dagstuhl dataset

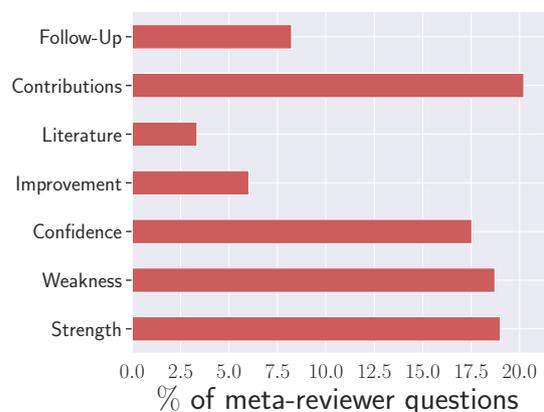


Figure 7: Distribution of meta-reviewer questions in our dataset.

### **A.2.1 Diversity of decisions in Dagstuhl dataset**

The Dagstuhl Dataset consists of 4 papers meta-reviewed by 4 senior academicians. The paper acceptance decisions is displayed in Figure 6. Across the 4 papers, only 1 paper had a unanimous reject decision, while all the other papers received diverse verdicts.

### **A.3 Prompts for all the experiments**

Our proposed method, ReMuSE, uses different prompting templates for the Initial Dialogue Generation, Feedback, and Refinement steps. The templates consist of an instruction specific to the particular ReMuSE step combined with our knowledge source, i.e., the paper’s reviews, title, and type (long or short). We use multiple prompts for the initial dialogue generation to account for model sensitivity. We name these prompts as follows: extensive prompt (simulates the study in [Balke et al. \(2023\)](#)), paraphrased prompt (paraphrase of the extensive prompt), and tl:dr prompt (3-sentence summary of the extensive prompt). We provide full templates for these prompts along with the ones used for Feedback and Refinement in Appendix §A.3.1, §A.3.2 and §A.3.3, respectively.

#### **A.3.1 Initial Dialogue Generation Prompts for all the experiments**

We show the different prompts for initial dialogue generation below:

## Extensive

## Prompt

### Extensive Prompt (Meta-Reviewing)

**Instruction:** Generate a multi-turn dialogue between a meta-reviewer and a dialogue agent for reviews about a paper. Every turn should alternate between "Dialogue Agent" and "Meta Reviewer" utterances. The meta-reviewer must explicitly tell whether the paper is accepted or rejected at the final turn. The dialogue ends with the meta-reviewer saying whether they accept or reject the paper. You need to simulate both the meta-reviewer and the dialogue agent. The dialogue agent has access to only the reviews and type of the paper. A meta-reviewer mainly asks questions and the dialogue agent should just answer from the reviews. A dialogue agent should never recommend anything/ give any opinions/ decide anything for the paper. A dialogue agent has no conscience of its own. A meta-reviewer can also ask follow-up questions and grill the dialogue agent for more information on the reviews. As a meta-reviewer, you should also weigh the importance of the condence of the reviewers while making a decision.

**Title:** [title]

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

### Extensive Prompt (Debates)

**Instruction:** Generate a multi-turn dialogue between a debate decision maker who needs to take a decision about which side wins a debate and a dialogue agent that has access to the arguments put forward by both the sides. Every turn should start with either "Dialogue Agent" followed by their utterance or "Decision Maker" followed by their utterance. You need to simulate both the decision maker and the dialogue agent. The decision maker does not have access to the arguments put forward by both the sides. The decision maker should just rely on the dialogue agent to know about the arguments from both the sides. A decision maker mainly asks questions and the dialogue agent answers. The dialogue agent should refrain from answering if its opinion/decision is sought. The dialogue agent has access to only the arguments of the debate. The dialogue agent should never decide on who wins the debate and should always take a neutral stand when any opinion is sought.

**Title:** "Abolish The Minimum Wage".

**Knowledge Source:** *For arguments: 'Let's talk about what would happen to real people in real families ...', Against Arguments: 'We need to abolish the minimum wage. And there's only one argument that matters on this issue, and it's the moral argument. Does the minimum wage make the world a better...'*

### Extensive Prompt (Product-Buying)

**Instruction:** Generate a multi-turn dialogue between a buyer who wants to buy a product and a dialogue agent for reviews about that product. Every turn should start with either "Dialogue Agent" followed by their utterance or "Buyer" followed by their utterance. The dialogue agent should always remain neutral and take a neutral stand in any case. The buyer should reach a decision to buy/not buy the paper at the final turn. You need to simulate both the buyer and the dialogue agent. The buyer does not have access to the reviews but the dialogue agent has access to the reviews and nothing else. It should answer any question about the reviews. A buyer mainly asks questions and the dialogue agent should refrain from offering advice if its opinion is sought. The dialogue agent should not recommend/advise anything regarding the product.

**Title:** "OP/TECH USA 1504372 Pro Loop Strap (Royal)".

**Knowledge Source:** *This strap is just what I needed for lugging my D700 + battery pack + It is comfortable and can handle a dslr and heavy lens. I bought this for the Nikon D700 camera and replaced the factory Nikon strap which is thinner. Very nice product with great adaptability and not made in China!...*

## Paraphrased

## Prompt

### Paraphrased Prompt (Meta-Reviewing)

**Instruction:** Create a multi-turn dialogue featuring a meta-reviewer and a dialogue agent discussing reviews of a paper. Each turn should begin with either "Dialogue Agent" or "Meta Reviewer" followed by their respective statements. The meta-reviewer must make a final decision to "accept" or "reject" the paper and explicitly state their choice at the end. The dialogue concludes with the meta-reviewer indicating whether they accept or reject the paper. The meta-reviewer, aware only of the paper's title, lacks access to reviews or information about the paper type and relies on the dialogue agent for this data. The dialogue agent has access solely to reviews and the paper type, responding neutrally and abstaining from offering opinions, recommendations, or decisions. If asked for its opinion, the dialogue agent must respond with "I do not know." The meta-reviewer primarily poses questions, and the dialogue agent answers based on the reviews, with the option to inquire about reviewer condence. The meta-reviewer may ask follow-up questions and press the dialogue agent for additional details from the reviews. It's important to note that the dialogue agent has no inherent conscience or personal opinions.

**Title:** [title]

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

### Paraphrased Prompt (Product-Buying)

**Instruction:** Create a multi-turn conversation between a prospective buyer interested in purchasing a specific product and a dialogue agent knowledgeable about reviews for that product. Each turn should commence with either "Dialogue Agent" followed by their statement or "Buyer" followed by theirs. The dialogue agent must maintain a neutral stance throughout and respond solely based on the reviews available. The buyer, who has not read the reviews, will make a definitive decision on whether to purchase the product or not in the final turn. The dialogue agent should refrain from expressing opinions and respond with 'I do not know' if asked for personal input. It's essential to clarify that the dialogue agent is not a seller of the product and possesses information solely from the reviews.

**Title:** "OP/TECH USA 1504372 Pro Loop Strap (Royal)".

**Knowledge Source:** *This strap is just what I needed for lugging my D700 + battery pack + It is comfortable and can handle a dslr and heavy lens. I bought this for the Nikon D700 camera and replaced the factory Nikon strap which is thinner. Very nice product with great adaptability and not made in China!...*

### Paraphrased Prompt (Debates)

**Instruction:** Craft a multi-turn conversation involving a decision maker tasked with determining the winning side in a debate, and a dialogue agent equipped with information on the arguments presented by both sides. Each turn should initiate with either "Dialogue Agent" or "Decision Maker" followed by their respective statements. The decision maker, who is unaware of the arguments, must make a conclusive decision on the debate's victor in the final turn. The dialogue agent should remain impartial, responding only based on the arguments, and stating "I do not know" if asked for its opinion. The decision maker relies solely on the dialogue agent for information about the arguments from both sides. The dialogue agent possesses access solely to the arguments and is prohibited from making decisions or expressing opinions on the debate outcome. **Title:** "Abolish The Minimum Wage".

**Knowledge Source:** *For arguments: 'Let's talk about what would happen to real people in real families if the U.S. were to go down this untested and quite radical path of eliminating the minimum wage as our opponents urge...', Against Arguments: 'We need to abolish the minimum wage. And there's only one argument that matters on this issue, and it's the moral argument. Does the minimum wage make the world a better place?...'*

### TL:DR

### Prompt

#### TL:DR Prompt for Initial Dialogue Generation (Meta-Reviewing)

**Instruction:** Generate a multi-turn dialogue between a meta-reviewer and a dialogue agent regarding a paper's reviews, with the meta-reviewer aiming to decide whether to accept or reject the paper. The dialogue agent, relying on reviews and aware of the paper type, must respond to the meta-reviewer's questions and refrain from providing opinions or recommendations, while the meta-reviewer weighs the importance of reviewer confidence in making the final decision.

**Title:** [title].

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

#### TL:DR Prompt for Initial Dialogue Generation (Product-Buying)

**Instruction:** Generate a multi-turn dialogue between a buyer and a neutral dialogue agent regarding a product, with the buyer making a final decision to buy or not. The dialogue agent, equipped with access to reviews, answers the buyer's questions, refraining from recommendations or opinions, while the buyer, without access to reviews, seeks information to reach a purchasing decision.

**Title:** "OP/TECH USA 1504372 Pro Loop Strap (Royal)".

**Knowledge Source:** *This strap is just what I needed for lugging my D700 + battery pack + It is comfortable and can handle a dslr and heavy lens. I bought this for the Nikon D700 camera and replaced the factory Nikon strap which is thinner. Very nice product with great adaptability and not made in China!...*

#### TL:DR Prompt for Initial Dialogue Generation (Debates)

**Instruction:** Simulate a multi-turn dialogue between a decision maker and a dialogue agent for a debate, where the decision maker, aware only of the debate's topic, seeks information from the dialogue agent, who has access to arguments from both sides. The decision maker aims to reach a conclusion about which side wins, while the dialogue agent remains neutral, providing answers solely based on the arguments presented, refraining from expressing opinions or making decisions.

**Title:** "Abolish The Minimum Wage".

**Knowledge Source:** *For arguments: 'Let's talk about what would happen to real people in real families if the U.S. were to go down this untested and quite radical path of eliminating the minimum wage as our opponents urge...', Against Arguments: 'We need to abolish the minimum wage. And there's only one argument that matters on this issue, and it's the moral argument. Does the minimum wage make the world a better place?...'*

### A.3.2 Prompt Templates for ReMuSE feedback

We have different versions of the same prompt used for multiple metrics. The prompt for refinement using all the metrics is provided below:

#### Prompt for ReMuSE feedback (Meta-Reviewing)

**Instruction:** Given the knowledge source and the dialogue, please provide actionable feedback to improve the dialogues. The feedback should just be for the overall dialogue and should start with 'Feedback:'. A Q2 F1 score, Q2 NLI score, KPrecision, and specificity scores follow each utterance. The Q2 and Kprecision scores measure the groundedness of the response to the provided knowledge source. The Specificity scores measure the technicality of the response. The feedback should try to improve all scores for the dialogue agent and the Specificity scores for the meta-reviewer.

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

**Dialogue:** *Meta-Reviewer: Hello Dialogue Agent. Can you tell me more about this paper?, F1: 0.0, NLI: 0.0, Kprec: 0.0, Specificity: 0.1 \n Dialogue Agent: Ofcourse! This is a paper about CodexGLUE, a machine learning benchmark dataset for code understanding and generation, F1: 0.12, NLI: 0.34, Kprec: 0.45, Specificity: 0.7 \n MetaReviewer: ...*

#### Prompt for ReMuSE feedback (Debates)

**Instruction:** Given the knowledge source and the dialogue, please provide actionable feedback to improve the dialogues. The feedback should just be for the overall dialogue and should start with 'Feedback:'. A Q2 F1 score, Q2 NLI score, KPrecision, and specificity scores follow each utterance. The Q2 and Kprecision scores measure the groundedness of the response to the provided knowledge source. The Specificity scores measure the technicality of the response. The feedback should try to improve all scores for the dialogue agent and the Specificity scores for the decision-maker.

**Title:** "Abolish The Minimum Wage".

**Knowledge Source:** *For arguments: 'Let's talk about what would happen to real people in real families if the U.S. were to go down this untested and quite radical path of eliminating the minimum wage as our opponents urge...', Against Arguments: 'We need to abolish the minimum wage. And there's only one argument that matters on this issue, and it's the moral argument. Does the minimum wage make the world a better place?...'*

**Dialogue:** *Decision Maker: Hello, can you tell the arguments for and against debating for minimum wages? F1: 0.2, NLI: 0.2, KPrec:0.01, Specificity: 0.2 \n Dialogue Agent: The debate over abolishing minimum wage centers on two perspectives... F1: 0.4, NLI: 0.39, KPrec: 0.45, Specificity: 0.6 ...*

#### Prompt for ReMuSE feedback (Product-Buying)

**Instruction:** Given the knowledge source and the dialogue, please provide actionable feedback to improve the dialogues. The feedback should just be for the overall dialogue and should start with 'Feedback:'. A Q2 F1 score, Q2 NLI score, KPrecision, and specificity scores follow each utterance. The Q2 and Kprecision scores measure the groundedness of the response to the provided knowledge source. The Specificity scores measure the technicality of the response. The feedback should try to improve all scores for the dialogue agent and the Specificity scores for the buyer.

**Knowledge Source:** *This strap is just what I needed for lugging my D700 + battery pack + It is comfortable and can handle a dslr and heavy lens. I bought this for the Nikon D700 camera and replaced the factory Nikon strap which is thinner..*

**Title:** "OP/TECH USA 1504372 Pro Loop Strap (Royal)".

**Dialogue:** *Buyer: Hi there, I am interested in buying the OP/TECH USA 1504372 Pro Loop Strap (Royal). can you tell me more about it? F1: 0.0, NLI: 0.0 \n Dialogue Agent: Sure! What would you like to know? F1: 0.0, NLI: 0.0, Kprec: 0.0, Spec: 0.1 \n Buyer: How comfortable are these for long travel hours? F1: 0.1, NLI: 0.1, Kprec: 0.05, Specificity: 0.4 \n Dialogue Agent: According to the reviews, its comfortable to wear for long hours, F1: 0.2, NLI: 0.2, Kprec: 0.25, Specificity: 0.5 \n Buyer: ...*

### A.3.3 Prompt Templates for ReMuSE refinement

We provide the refinement prompts here:

#### Prompt for Refinement (Product-Buying)

**Instruction:** Given the feedback, knowledge source, and dialogue, improve the dialogue based on the feedback provided. The output should just be the new dialogue.

**Knowledge Source:** *This strap is just what I needed for lugging my D700 + battery pack + It is comfortable and can handle a dslr and heavy lens. I bought this for the Nikon D700 camera and replaced the factory Nikon strap which is thinner. Very nice product with great adaptability and not made in China!...*

**Feedback:** *Here are some specific feedback points: 1. The buyer is hesitant about the price, and the dialogue agent could provide more information about the product's value and benefits to address these concerns. 2. The buyer asks about the reviews and ratings, and the dialogue agent could provide more specific examples or testimonials to showcase the product's positive feedback. 3. The dialogue agent could provide more information about the return policy and the retailer's customer service to alleviate any concerns the buyer may have.*

**Dialogue:** *Buyer: Hi there, I am interested in buying the OP/TECH USA 1504372 Pro Loop Strap (Royal). can you tell me more about it? \n Dialogue Agent: Sure! What would you like to know? \n Buyer: How comfortable are these for long travel hours? \n Dialogue Agent: According to the reviews, its comfortable to wear for long hours, \n Buyer: ...*

#### Prompt for Refinement (Meta-Reviewing)

**Instruction:** Given the feedback, knowledge source, and dialogue improve the dialogue based on the feedback provided. The output should just be the new dialogue.

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

**Feedback:** *It would be helpful to provide a clearer explanation of the tasks' relevance to software development in response to reviewer one's concerns. Additionally, it may be useful to mention that the paper discusses the effectiveness of the baselines and provides an ablation study. Finally, it would be good to acknowledge reviewer three's comment explicitly and mention any plans to revise the descriptions of the datasets.*

**Dialogue:** *Meta-Reviewer: Hello Dialogue Agent. Can you tell me more about this paper? \n Dialogue Agent: Ofcourse! This is a paper about CodexGLUE, a machine learning benchmark dataset for code understanding and generation, \n MetaReviewer: Interesting. Can you give me some more information about the reviews that have been given for this paper? \n Dialogue Agent: ...*

Domain	Model	Rewards	K-Prec	Q2-F1	Q2-NLI	Spec.
Meta-Reviewing	ChatGPT	-	0.387 ± 0.012	0.147 ± 0.011	0.152 ± 0.011	0.451 ± 0.127
	LLaMa	-	0.341 ± 0.014	0.170 ± 0.130	0.170 ± 0.014	0.461 ± 0.081
	Mistral	-	0.377 ± 0.012	0.151 ± 0.022	0.161 ± 0.014	0.407 ± 0.015
	Mixtral	-	0.378 ± 0.014	0.171 ± 0.014	0.188 ± 0.021	0.422 ± 0.021
	ChatGPT	K-Prec	0.679 ± 0.022	0.177 ± 0.121	0.197 ± 0.025	0.552 ± 0.034
	LLaMa	K-Prec	0.633 ± 0.016	0.162 ± 0.141	0.163 ± 0.012	0.563 ± 0.021
	Mistral	K-Prec	0.681 ± 0.014	0.242 ± 0.014	0.231 ± 0.021	0.611 ± 0.024
	Mixtral	K-Prec	0.702 ± 0.011	0.268 ± 0.021	0.263 ± 0.015	0.646 ± 0.026
	ChatGPT	Q2 (F1 + NLI)	0.689 ± 0.013	0.171 ± 0.011	0.173 ± 0.013	0.648 ± 0.018
	LLaMa	Q2 (F1 + NLI)	0.646 ± 0.017	0.178 ± 0.029	0.189 ± 0.013	0.571 ± 0.026
	Mistral	Q2 (F1 + NLI)	0.706 ± 0.017	0.268 ± 0.034	0.272 ± 0.012	0.626 ± 0.013
	Mixtral	Q2 (F1 + NLI)	0.721 ± 0.014	0.271 ± 0.016	0.265 ± 0.012	0.701 ± 0.028
	ChatGPT	Spec.	0.692 ± 0.015	0.180 ± 0.023	0.181 ± 0.019	0.652 ± 0.032
	LLaMa	Spec.	0.681 ± 0.010	0.183 ± 0.021	0.191 ± 0.016	0.592 ± 0.035
	Mistral	Spec.	0.703 ± 0.021	0.292 ± 0.021	0.300 ± 0.021	0.639 ± 0.020
	Mixtral	Spec.	0.704 ± 0.012	0.285 ± 0.017	0.269 ± 0.014	0.689 ± 0.020
	ChatGPT	K-Prec + Q2	0.701 ± 0.021	0.178 ± 0.014	0.188 ± 0.067	0.648 ± 0.024
	LLaMa	K-Prec + Q2	0.712 ± 0.023	0.177 ± 0.028	0.189 ± 0.014	0.555 ± 0.038
	Mistral	K-Prec + Q2	0.691 ± 0.013	0.288 ± 0.015	0.290 ± 0.186	0.611 ± 0.015
	Mixtral	K-Prec + Q2	0.731 ± 0.032	0.276 ± 0.189	0.291 ± 0.178	0.672 ± 0.112
	ChatGPT	Spec. + K-Prec	0.711 ± 0.031	0.184 ± 0.012	0.182 ± 0.131	0.665 ± 0.021
	LLaMa	Spec. + K-Prec	0.704 ± 0.067	0.184 ± 0.034	0.187 ± 0.024	0.578 ± 0.024
	Mistral	Spec. + K-Prec	0.686 ± 0.021	0.281 ± 0.071	0.299 ± 0.045	0.591 ± 0.021
	Mixtral	Spec. + K-Prec	0.728 ± 0.021	0.282 ± 0.031	0.321 ± 0.141	0.711 ± 0.027
	ChatGPT	Spec. + Q2	0.721 ± 0.172	0.167 ± 0.021	0.174 ± 0.019	0.661 ± 0.027
	LLaMa	Spec. + Q2	0.718 ± 0.089	0.174 ± 0.056	0.188 ± 0.065	0.588 ± 0.041
	Mistral	Spec. + Q2	0.688 ± 0.129	0.267 ± 0.014	0.271 ± 0.011	0.578 ± 0.018
	Mixtral	Spec. + Q2	0.741 ± 0.121	0.282 ± 0.073	0.330 ± 0.021	0.676 ± 0.043
	ChatGPT	K-Prec + Q2 + Spec.	0.762 ± 0.014	0.241 ± 0.021	0.241 ± 0.178	0.614 ± 0.172
	LLaMa	K-Prec + Q2 + Spec.	0.740 ± 0.021	0.227 ± 0.018	0.238 ± 0.019	0.626 ± 0.016
	Mistral	K-Prec + Q2 + Spec.	0.762 ± 0.013	0.304 ± 0.016	0.326 ± 0.018	0.646 ± 0.221
	Mixtral	K-Prec + Q2 + Spec.	<b>0.781 ± 0.012</b>	<b>0.322 ± 0.012</b>	<b>0.344 ± 0.017</b>	<b>0.721 ± 0.018</b>

Table 7: Performance of ReMuSE with different combinations of rewards on the Meta-Reviewing Domain.

Domain	Model	Rewards	K-Prec	Q2-F1	Q2-NLI	Spec.
Product Buying	ChatGPT	-	0.497 ± 0.023	0.202 ± 0.014	0.203 ± 0.011	0.311 ± 0.024
	LLaMa	-	0.462 ± 0.022	0.203 ± 0.034	0.206 ± 0.044	0.448 ± 0.017
	Mistral	-	0.523 ± 0.015	0.225 ± 0.022	0.225 ± 0.019	0.427 ± 0.028
	Mixtral	-	0.461 ± 0.026	0.263 ± 0.039	0.323 ± 0.045	0.424 ± 0.076
	ChatGPT	K-Prec	0.539 ± 0.010	0.226 ± 0.012	0.240 ± 0.024	0.441 ± 0.012
	LLaMa	K-Prec	0.580 ± 0.032	0.236 ± 0.014	0.246 ± 0.321	0.541 ± 0.032
	Mistral	K-Prec	0.582 ± 0.036	0.267 ± 0.153	0.267 ± 0.242	0.475 ± 0.128
	Mixtral	K-Prec	0.622 ± 0.010	0.289 ± 0.011	0.287 ± 0.152	0.529 ± 0.009
	ChatGPT	Q2 (F1 + NLI)	0.533 ± 0.029	0.287 ± 0.018	0.294 ± 0.015	0.438 ± 0.015
	LLaMa	Q2 (F1 + NLI)	0.544 ± 0.006	0.242 ± 0.018	0.257 ± 0.125	0.522 ± 0.025
	Mistral	Q2 (F1 + NLI)	0.593 ± 0.032	0.278 ± 0.036	0.282 ± 0.036	0.467 ± 0.042
	Mixtral	Q2 (F1 + NLI)	0.662 ± 0.004	0.335 ± 0.013	0.347 ± 0.011	0.525 ± 0.004
	ChatGPT	Spec.	0.643 ± 0.012	0.267 ± 0.019	0.276 ± 0.019	0.453 ± 0.039
	LLaMa	Spec.	0.621 ± 0.023	0.226 ± 0.027	0.239 ± 0.030	0.589 ± 0.029
	Mistral	Spec.	0.634 ± 0.037	0.258 ± 0.054	0.276 ± 0.069	0.602 ± 0.063
	Mixtral	Spec.	0.642 ± 0.015	0.342 ± 0.032	0.344 ± 0.035	0.592 ± 0.001
	ChatGPT	SR	0.510 ± 0.018	0.215 ± 0.020	0.220 ± 0.018	0.330 ± 0.030
	LLaMa	SR	0.485 ± 0.020	0.218 ± 0.022	0.225 ± 0.022	0.460 ± 0.018
	Mistral	SR	0.540 ± 0.022	0.240 ± 0.020	0.245 ± 0.019	0.440 ± 0.025
	Mixtral	SR	0.495 ± 0.025	0.270 ± 0.030	0.310 ± 0.032	0.435 ± 0.040
	ChatGPT	SC+SR	0.520 ± 0.015	0.225 ± 0.025	0.230 ± 0.021	0.345 ± 0.032
	LLaMa	SC+SR	0.495 ± 0.018	0.230 ± 0.020	0.235 ± 0.018	0.470 ± 0.020
	Mistral	SC+SR	0.550 ± 0.020	0.250 ± 0.025	0.260 ± 0.020	0.455 ± 0.030
	Mixtral	SC+SR	0.505 ± 0.022	0.275 ± 0.028	0.315 ± 0.030	0.445 ± 0.035
	ChatGPT	K-Prec + Q2	0.623 ± 0.017	0.243 ± 0.035	0.260 ± 0.121	0.463 ± 0.051
	LLaMa	K-Prec + Q2	0.631 ± 0.046	0.234 ± 0.129	0.237 ± 0.038	0.552 ± 0.042
	Mistral	K-Prec + Q2	0.641 ± 0.043	0.242 ± 0.117	0.264 ± 0.015	0.588 ± 0.052
	Mixtral	K-Prec + Q2	0.671 ± 0.432	0.326 ± 0.012	0.341 ± 0.028	0.602 ± 0.016
	ChatGPT	Spec. + K-Prec	0.612 ± 0.015	0.241 ± 0.016	0.311 ± 0.178	0.524 ± 0.028
	LLaMa	Spec. + K-Prec	0.582 ± 0.016	0.248 ± 0.018	0.254 ± 0.048	0.581 ± 0.063
	Mistral	Spec. + K-Prec	0.633 ± 0.061	0.241 ± 0.017	0.263 ± 0.018	0.572 ± 0.061
	Mixtral	Spec. + K-Prec	0.664 ± 0.038	0.327 ± 0.019	0.341 ± 0.026	0.585 ± 0.042
	ChatGPT	Spec. + Q2	0.652 ± 0.041	0.267 ± 0.045	0.287 ± 0.078	0.535 ± 0.056
	LLaMa	Spec. + Q2	0.621 ± 0.064	0.257 ± 0.054	0.268 ± 0.068	0.591 ± 0.017
	Mistral	Spec. + Q2	0.641 ± 0.074	0.256 ± 0.076	0.322 ± 0.059	0.622 ± 0.078
	Mixtral	Spec. + Q2	0.676 ± 0.051	0.358 ± 0.067	0.361 ± 0.071	0.621 ± 0.051
	ChatGPT	K-Prec + Q2 + Spec.	0.671 ± 0.014	0.284 ± 0.151	0.314 ± 0.115	0.555 ± 0.157
	LLaMa	K-Prec + Q2 + Spec.	0.637 ± 0.015	0.286 ± 0.131	0.299 ± 0.141	0.581 ± 0.231
	Mistral	K-Prec + Q2 + Spec.	0.697 ± 0.051	0.339 ± 0.182	0.351 ± 0.132	0.604 ± 0.214
	Mixtral	K-Prec + Q2 + Spec.	<b>0.713 ± 0.141</b>	<b>0.372 ± 0.121</b>	<b>0.404 ± 0.151</b>	<b>0.652 ± 0.191</b>

Table 8: Performance of ReMuSE on the Product Buying Domain. Best results are indicated in **bold**.

Domain	Model	Rewards	K-Prec	Q2-F1	Q2-NLI	Spec.
Debates	ChatGPT	-	0.374 ± 0.011	0.171 ± 0.121	0.202 ± 0.141	0.398 ± 0.141
	LLaMa	-	0.377 ± 0.021	0.177 ± 0.214	0.209 ± 0.151	0.427 ± 0.115
	Mistral	-	0.406 ± 0.022	0.223 ± 0.113	0.246 ± 0.133	0.443 ± 0.122
	Mixtral	-	0.411 ± 0.015	0.208 ± 0.141	0.226 ± 0.115	0.451 ± 0.114
	ChatGPT	SR	0.390 ± 0.012	0.180 ± 0.120	0.215 ± 0.130	0.420 ± 0.120
	LLaMa	SR	0.395 ± 0.015	0.185 ± 0.125	0.220 ± 0.135	0.430 ± 0.110
	Mistral	SR	0.420 ± 0.018	0.230 ± 0.110	0.245 ± 0.120	0.450 ± 0.115
	Mixtral	SR	0.430 ± 0.016	0.235 ± 0.115	0.250 ± 0.118	0.460 ± 0.110
	ChatGPT	SC+SR	0.400 ± 0.010	0.190 ± 0.110	0.225 ± 0.125	0.440 ± 0.105
	LLaMa	SC+SR	0.405 ± 0.012	0.195 ± 0.112	0.230 ± 0.120	0.450 ± 0.100
	Mistral	SC+SR	0.430 ± 0.015	0.240 ± 0.108	0.255 ± 0.115	0.470 ± 0.108
	Mixtral	SC+SR	0.440 ± 0.014	0.245 ± 0.110	0.260 ± 0.112	0.480 ± 0.105
	ChatGPT	K-Prec	0.444 ± 0.006	0.180 ± 0.131	0.215 ± 0.151	0.482 ± 0.221
	LLaMa	K-Prec	0.414 ± 0.012	0.187 ± 0.152	0.214 ± 0.042	0.529 ± 0.033
	Mistral	K-Prec	0.486 ± 0.004	0.235 ± 0.151	0.244 ± 0.041	0.547 ± 0.003
	Mixtral	K-Prec	0.483 ± 0.009	0.246 ± 0.015	0.241 ± 0.141	0.581 ± 0.033
	ChatGPT	Q2 (F1 + NLI)	0.421 ± 0.028	0.194 ± 0.020	0.228 ± 0.024	0.481 ± 0.027
	LLaMa	Q2 (F1 + NLI)	0.412 ± 0.027	0.216 ± 0.039	0.224 ± 0.045	0.518 ± 0.033
	Mistral	Q2 (F1 + NLI)	0.487 ± 0.039	0.251 ± 0.043	0.277 ± 0.045	0.552 ± 0.015
	Mixtral	Q2 (F1 + NLI)	0.504 ± 0.009	0.268 ± 0.034	0.288 ± 0.010	0.574 ± 0.022
	ChatGPT	Spec.	0.414 ± 0.017	0.177 ± 0.018	0.186 ± 0.027	0.494 ± 0.018
	LLaMa	Spec.	0.400 ± 0.031	0.181 ± 0.047	0.190 ± 0.044	0.540 ± 0.032
	Mistral	Spec.	0.484 ± 0.024	0.249 ± 0.046	0.250 ± 0.048	0.581 ± 0.018
	Mixtral	Spec.	0.491 ± 0.019	0.252 ± 0.012	0.261 ± 0.017	0.601 ± 0.010
	ChatGPT	K-Prec + Q2	0.419 ± 0.057	0.208 ± 0.054	0.214 ± 0.043	0.526 ± 0.051
	LLaMa	K-Prec + Q2	0.411 ± 0.064	0.192 ± 0.051	0.201 ± 0.054	0.537 ± 0.018
	Mistral	K-Prec + Q2	0.473 ± 0.051	0.262 ± 0.043	0.278 ± 0.052	0.574 ± 0.087
	Mixtral	K-Prec + Q2	0.494 ± 0.064	0.268 ± 0.018	0.267 ± 0.064	0.628 ± 0.065
	ChatGPT	Spec. + K-Prec	0.422 ± 0.052	0.201 ± 0.044	0.209 ± 0.116	0.484 ± 0.054
	LLaMa	Spec. + K-Prec	0.428 ± 0.089	0.194 ± 0.062	0.205 ± 0.067	0.542 ± 0.144
	Mistral	Spec. + K-Prec	0.489 ± 0.057	0.244 ± 0.067	0.259 ± 0.089	0.584 ± 0.056
	Mixtral	Spec. + K-Prec	0.484 ± 0.078	0.252 ± 0.056	0.252 ± 0.018	0.624 ± 0.076
	ChatGPT	Spec. + Q2	0.448 ± 0.197	0.201 ± 0.069	0.212 ± 0.068	0.502 ± 0.074
	LLaMa	Spec. + Q2	0.421 ± 0.079	0.221 ± 0.168	0.248 ± 0.156	0.551 ± 0.018
	Mistral	Spec. + Q2	0.486 ± 0.058	0.248 ± 0.188	0.268 ± 0.076	0.591 ± 0.054
	Mixtral	Spec. + Q2	0.510 ± 0.098	0.278 ± 0.088	0.281 ± 0.045	0.621 ± 0.077
	ChatGPT	K-Prec + Q2 + Spec.	0.501 ± 0.016	0.221 ± 0.114	0.232 ± 0.168	0.535 ± 0.132
	LLaMa	K-Prec + Q2 + Spec.	0.441 ± 0.131	0.234 ± 0.178	0.244 ± 0.132	0.565 ± 0.151
	Mistral	K-Prec + Q2 + Spec.	0.501 ± 0.112	0.274 ± 0.213	0.291 ± 0.121	0.627 ± 0.189
	Mixtral	K-Prec + Q2 + Spec.	<b>0.522 ± 0.131</b>	<b>0.277 ± 0.122</b>	<b>0.298 ± 0.152</b>	<b>0.628 ± 0.121</b>

Table 9: Performance of ReMuSE on the Debates Domain. Best results are indicated in **bold**.

### Prompt for Refinement (Debates)

**Instruction:** Given the feedback, knowledge source, and dialogue improve the dialogue based on the feedback provided. The output should just be the new dialogue.

**Knowledge Source: For arguments:** *'Let's talk about what would happen to real people in real families if the U.S. were to go down this untested and quite radical path of eliminating the minimum wage as our opponents urge...'*, **Against Arguments:** *'We need to abolish the minimum wage. And there's only one argument that matters on this issue, and it's the moral argument. Does the minimum wage make the world a better place?...'*

**Feedback:** *It could benefit from more specific and concise answers to the questions asked by the decision-maker. 1. The Q2 F1 score could be improved by providing more specific examples and evidence to support the arguments presented. For example, when discussing the argument for abolishing the minimum wage, the dialogue agent could provide examples of how higher labor costs have negatively impacted businesses in the past. 2. The Q2 NLI score could be improved by ensuring that the dialogue agent's responses are coherent and make sense in the context of the conversation. For example, when discussing the argument for keeping the minimum wage, the dialogue agent could provide more specific examples of how a living wage would benefit low-income families.*

**Dialogue:** *Decision Maker: Hello, can you tell the arguments for and against debating for minimum wages? \n Dialogue Agent: The debate over abolishing the minimum wage centers on two perspectives. Supporters claim it harms businesses and argue for market-driven wages without government interference. Opponents assert the minimum wage is crucial to shield low-income workers from exploitation, contending that without it, wages might be insufficient for basic needs. \n Decision Maker: ..*

### Prompt for Generic feedback (Meta-Reviewing)

**Instruction:** Given the knowledge source and the dialogue, please provide **feedback** to improve the dialogues. The feedback should just be for the overall dialogue and should start with 'Feedback:'.

**Knowledge Source: Review 1:** *The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ...*, **Review 2:** *...*, **Review 3:** *...*

**Dialogue:** *Meta-Reviewer: Hello Dialogue Agent. Can you tell me more about this paper? \n Dialogue Agent: Ofcourse! This is a paper about CodexGLUE, a machine learning benchmark dataset for code understanding and generation, \n MetaReviewer: ...*

### Prompt for Actionable feedback (Meta-Reviewing)

**Instruction:** Given the knowledge source and the dialogue, please provide **actionable feedback** to improve the dialogues. The feedback should be for the overall dialogue and start with 'Feedback:.' The feedback should improve the groundedness of the dialogue agent's utterances to the knowledge source. The feedback should also increase the specificity (higher technical details) of utterances for the dialogue agent and the meta-reviewer.

**Knowledge Source: Review 1:** *The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ...*, **Review 2:** *...*, **Review 3:** *...*

**Dialogue:** *Meta-Reviewer: Hello Dialogue Agent. Can you tell me more about this paper? \n Dialogue Agent: Ofcourse! This is a paper about CodexGLUE, a machine learning benchmark dataset for code understanding and generation, \n MetaReviewer: ...*

## A.4 Cross-Domain analysis for the performance of ReMuSE

Table 7 shows the detailed performance of ReMuSE with multiple combinations of rewards for the meta-reviewing domain. We discuss the cross-domain analysis of ReMuSE on two other decision-making domains namely debates and product reviews below.

### A.4.1 Domains

In order to test the robustness of our approach, ReMuSE, we experiment with two other domains of decision-making based on lengthy input documents. In *product-buying*, a buyer needs to decide for or against a product based on the provided customer reviews. In *debates*, a listener must decide on one of the stances based on the provided arguments.

In this section, we initially discuss the datasets we employed to perform the analysis, followed by a discussion on our experimental findings.

### A.4.2 Data

For *product-buying*, we use the Helpful Reviews Dataset (Gamzu et al., 2021). The reviews are available for 123 products, and each review comes annotated with a helpfulness score. Following the recommendation of Gamzu et al. (2021), we consider all the reviews scored above 1, which are the most helpful for making buying decisions. For *debates*, we use the IQ2 Debates Dataset (Zhang et al., 2016). The dataset consists of Oxford-Style Debates on 108 topics. Each debate is a dialogue between multiple parties namely, moderator, pro-side, and con-side. For ease of use, we only con-

sider the opening arguments from the pro- and con-sides of the debate for each topic.

### A.4.3 Results

In the case of product buying, we use all the helpful review sentences available for a product as the knowledge source, whereas the arguments from pro and con-sides of a topic in case of debates. We use similar prompts as in Sec §A.3 but simulate different scenarios for both domains and then apply our method as is. The detailed prompts for initial dialogue generation, feedback, and refinement for both the domains are provided in §A.3.1, §A.3.2 and §A.3.3 respectively. We report the results for the *product-buying* and *debates* domain in Tables 8 and 9 in §A.4 respectively. The best results are attained by using an amalgamation of all the refinement metrics, similar to the meta-reviewing domain. We observe a similar trend in terms of the best-performing model as well where the Mixtral model excels across all the domains.

## A.5 Ablation Study

### A.5.1 Single vs Joint Reward Optimization

We show the scores for single vs joint score optimization in Table 10. We observe massive improvements when using the amalgamation of all the refinement scores.

Reward	K-Prec	Q2-F1	Q2-NL	Spec.
K-Prec	0.70 <sub>.01</sub>	0.27 <sub>.02</sub>	0.26 <sub>.02</sub>	0.65 <sub>.03</sub>
Q2	0.72 <sub>.01</sub>	0.27 <sub>.02</sub>	0.27 <sub>.01</sub>	0.70 <sub>.03</sub>
Spec.	0.70 <sub>.01</sub>	0.29 <sub>.02</sub>	0.27 <sub>.01</sub>	0.69 <sub>.02</sub>
K-Prec + Q2	0.73 <sub>.03</sub>	0.28 <sub>.15</sub>	0.29 <sub>.18</sub>	0.67 <sub>.11</sub>
Spec. + K-Prec	0.73 <sub>.02</sub>	0.28 <sub>.03</sub>	0.32 <sub>.14</sub>	0.71 <sub>.03</sub>
Spec. + Q2	0.74 <sub>.12</sub>	0.28 <sub>.07</sub>	0.33 <sub>.02</sub>	0.68 <sub>.04</sub>
ReMuSE	<b>0.78<sub>.01</sub></b>	<b>0.32<sub>.01</sub></b>	<b>0.34<sub>.02</sub></b>	<b>0.72<sub>.02</sub></b>

Table 10: Single vs Joint-Score Optimization of ReMuSE for our best-performing model, **Mixtral** in terms of various refinement metrics.

### A.5.2 Multi-Iterative Refinement

In Fig 8, we plot the scores for each iteration of ReMuSE up to 10 epochs and observe that the scores plateau after the first epoch, indicating superior feedback quality in the first round, consistent with prior work (Madaan et al., 2023).

### A.5.3 Generic vs Actionable vs Multi-Reward Feedback

We plot the various metrics for these multiple feedback strategies in Fig 9. We observe that the feedback-prompting in ReMuSE leads to the maximum improvements. We also show a few examples

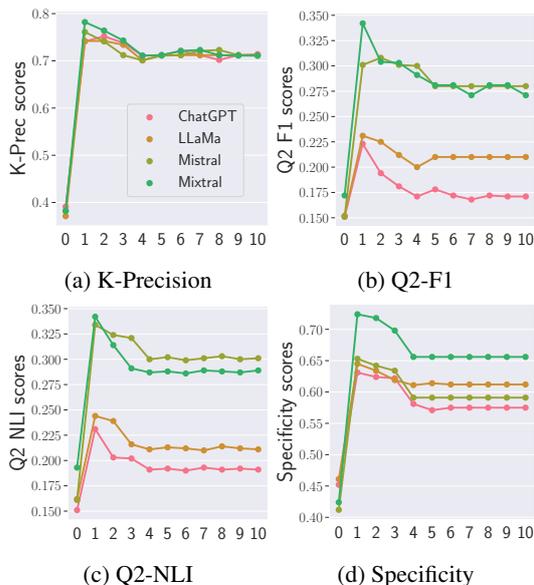


Figure 8: For the iterative refinement version of ReMuSE, we plot the (a) K-Precision, (b) Q2-F1, (c) Q2-NLI, and, (d) Specificity scores with varying iterations (0 – 4). All the results are averaged over 3 prompts.

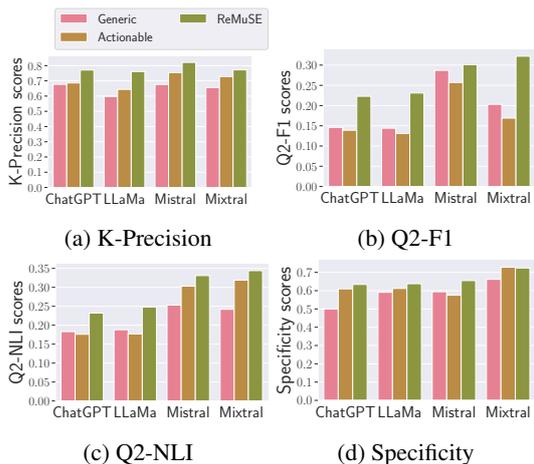


Figure 9: Performances for the different feedback types Generic, Actionable, and ReMuSE. We plot the (a) K-Precision, (b) Q2-F1, (c) Q2-NLI, and, (d) Specificity scores for each (averaged over 3 prompts).

of the feedback generated for a dialogue using the prompts discussed above and our proposed method, ReMuSE in Table 11.

### A.5.4 Correlation of human evaluation metrics

We plot the correlation of the different metrics used for evaluating the dialogue agent’s responses, namely, fluency, faithfulness, relevance, objectivity, and helpfulness, in Fig 5. We observe a strong correlation between faithfulness, objectiv-

<b>Strategy</b>	<b>Feedback</b>
Generic	The agent mostly summarized the concerns raised by the reviewers without suggesting any specific actions that could be taken to address those concerns. In the future, the agent could provide more specific suggestions or questions to help the meta-reviewer make a more informed decision. Additionally, the agent could try to paraphrase or restate the concerns raised by the reviewers in a clearer or more concise way to facilitate the conversation.
Actionable	Overall, it is advisable for the authors to further clarify their proposed algorithm by addressing the issues raised by the reviewers. For instance, they can provide a better definition of transportation polytope in Eq. (2) and explain the process of generating two views of the input images in more detail. It is also recommended that they compare their method with the state-of-the-art consensus clustering methods that are not deep learning-based as suggested by Reviewer 1. To address the concerns over the quality of the experiments, the authors can provide a more thorough experimental investigation and comparison, as well as a discussion of the novelty and contributions of their method with respect to similar methods proposed in prior work.
ReMuSE	Here are some suggestions to improve the dialogue: 1. Provide a more detailed summary of the reviewers' concerns, such as the specific areas where the paper is lacking in novelty or clarity. For example, the dialogue could mention the reviewers' comments about the outdated literature, vague algorithm description, and lack of analysis of results. 2. Offer specific recommendations to address the reviewers' concerns. For instance, the dialogue could suggest that the authors update the literature review with more recent work on consensus clustering, provide a more detailed explanation of the algorithm, and perform additional analysis of the results. 3. Highlight the positive aspects of the paper, such as the novelty of combining ensemble methods with deep clustering. The dialogue could also mention the reviewers' comments about the potential effectiveness of the proposed ensemble deep clustering algorithm.

Table 11: Different types of feedback generated for the same dialogue using multiple feedback prompting strategies.

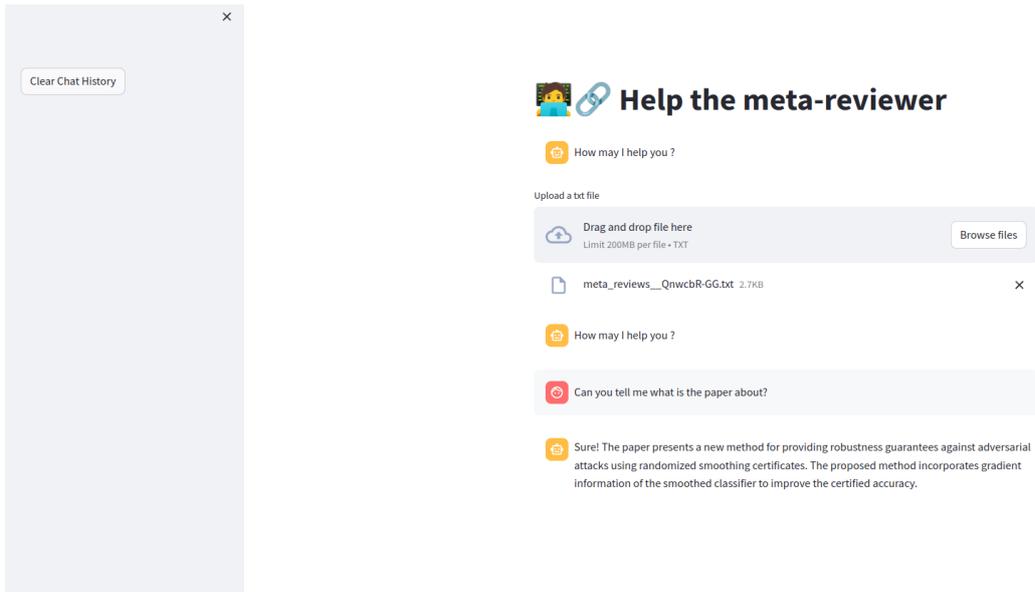


Figure 10: Screenshot of the Interface for assisting meta-reviewers with a dialogue agent

ity, and helpfulness. This further indicates a low tolerance for unfaithful responses in the decision-making process. The models finetuned on our data generate more faithful responses as demonstrated through the scores in human evaluation in Table 5 in Sec §4.2.

### A.6 Resulting Dataset

As one of the contributions of this work, we make the synthetically generated dataset for the best-performing ReMuSE configuration publicly available. We release 3,064 meta-reviewing, 128 product-buying, and 108 debate dialogues, respectively. We display the basic statistics for our resulting dataset in Table 12. Meta-Reviewing dialogues show the most diversity (unigram vocabulary size) in agent and meta-reviewer utterances, reflecting the specific terminology associated with the peer-reviewing domain. The debates domain closely follows due to diverse topics, while the product-buying domain exhibits the least diversity, reflecting the domain’s simplicity. Meta-Reviewing dialogues are notably longer, with more turns depicting the longer reasoning steps in expert tasks. The turn distributions in our dataset are highlighted in Figures 11, 12, 13 for meta-reviewing, product buying, and debates datasets, respectively.

### A.7 Hyperparameters for Dialogue Response Generation

All models are fine-tuned for 10 epochs. We perform a grid search on learning rates  $\{1e-4, 3e-$

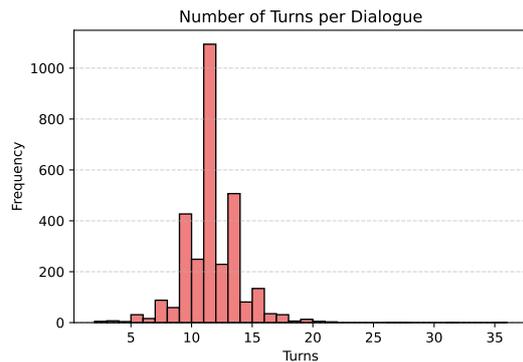


Figure 11: Distribution of Number of turns in Meta-Reviewing Dataset

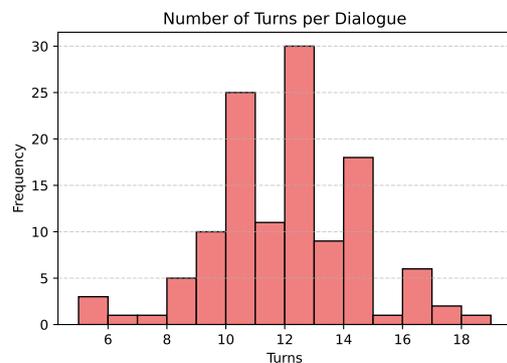


Figure 12: Distribution of Number of turns in Product Buying Dataset

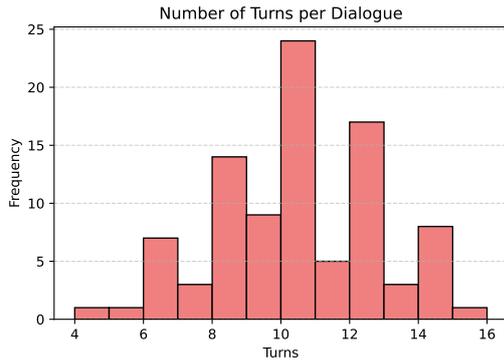


Figure 13: Distribution of Number of turns in Debates Dataset

4, 5e - 5, 6e - 5} and select the best model based on the validation set performance. The results are averaged over 3 random seeds. For ReMuSE, we use the best-performing model, Mixtral, and update the responses of ReMuSE for 1 iteration, given our prior findings. We adapt the same prompts as in A.3 for the response generation task.

#### Response generation Prompt (Meta-Reviewing)

**Instruction:** Generate a response corresponding to a dialogue for a conversation between a meta-reviewer and a dialogue agent for reviews about a paper. The dialogue agent has access to the dialogue history, reviews, and the type of paper. The dialogue agent should just answer from the reviews. A dialogue agent should never recommend anything/ give any opinions/ decide anything for the paper. A dialogue agent has no conscience of its own.

**Title:** [title]

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...* **Dialogue History:** *Meta-reviewer: Hello, dialogue agent. Thank you for joining me today to discuss the reviews for the paper [title]. Can you please start by summarizing the main contributions of the paper and how it addresses the problem it sets out to solve?*

#### Prompt for Response generation Re-MuSE feedback

**Instruction:** Given the knowledge source, dialogue history and the response, please provide actionable feedback to improve the response. The feedback should be for the response and should start with 'Feedback:'. A Q2 F1 score, Q2 NLI score, KPrecision, and specificity scores. The Q2 and Kprecision scores measure the groundedness of the response to the provided knowledge source. The Specificity scores measure the technicality of the response. The feedback should try to improve all scores for the dialogue agent

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

**Dialogue History:** *Meta-reviewer: Hello, dialogue agent. Thank you for joining me today to discuss the reviews for the paper [title]. Can you please start by summarizing the main contributions of the paper and how it addresses the problem it sets out to solve?*

**Response:** *Dialogue Agent: Sure. The paper proposes a method based on importance and coherence to evaluate whether a partitioning of a neural network exhibits modular characteristics. F1:0.12, NLI:0.14 and Specificity: 0.1*

#### Prompt for Response generation Re-MuSE Refinement

**Instruction:** Given the feedback, knowledge source, dialogue history and the previous response, improve the response based on the feedback provided. The output should just be the new dialogue.

**Knowledge Source:** *Review 1: The paper proposes a benchmark to measure the quality of the systems in the growing domain of code intelligence ..., Review 2: ..., Review 3: ...*

**Feedback:** *It would be helpful to provide a clearer explanation of the tasks' relevance to software development in response to reviewer one's concerns. Additionally, it may be useful to mention that the paper discusses the effectiveness of the baselines and provides an ablation study. Finally, it would be good to acknowledge reviewer three's comment explicitly and mention any plans to revise the descriptions of the datasets.*

**Dialogue history:** *Meta-reviewer: Hello, dialogue agent. Thank you for joining me today to discuss the reviews for the paper [title]. Can you please start by summarizing the main contributions of the paper and how it addresses the problem it sets out to solve?*

**Previous Response:** *Dialogue Agent: Sure. The paper proposes a method based on importance and coherence to evaluate whether a partitioning of a neural network exhibits modular characteristics.*

## A.8 Meta-Reviewer Question Distribution

To quantify the different types of questions, we quantitatively analyze the different meta-reviewer questions in our dataset. To come up with the categories, we use the 'Dagstuhl' dataset and annotate the questions in that dataset. We find that most of the meta-reviewer questions are related to Contributions, Strengths, Weaknesses, Confidence, Avenues for Improvement, and Literature Survey and Follow-Up. We then prompt GPT 3.5-turbo to provide us with these labels for all the questions in our dataset. The distribution of questions in our dataset is provided in Fig 7. We observed that questions related to contributions, strengths, weaknesses, contributions, and confidence (~ 20%) occur across all dialogues. These questions are the most important aspects for consideration while writing a meta-review, which is in line with the findings of Hossain et al. (2025). Thus, the questions in our dataset can serve as an effective resource for serving as a proxy meta-reviewer while training dialogue agents.

## A.9 Error Analysis of the dialogue agent responses

Following Li et al. (2023b), we quantitatively analyze a set of 50 dialogues for error types in the dialogue agent's responses. Out of 150 dialogue agent turns, we observed 50 responses that may be considered erroneous. We identified four different kinds of errors in the responses: 1) **Irrelevant answers:** 75% errors are incurred when the dialogue agent tries to answer questions that are based on external knowledge. For instance, for a meta-reviewer question on the literature survey of the

Domain	# dial.	# DA. tok. (av.)	# Seek. tok. (av.)	# tur.
Meta-Rev.	3,064	54.66	30.95	12.24
Product-Buy.	123	28.87	13.49	11.55
Debates	108	48.51	18.16	10.0

Table 12: Basic Statistics for the generated dataset. The table shows the number of dialogues, the average number of tokens for the dialogue agent and seeker utterances, and the number of turns.

paper, the dialogue agent answers about the novelty of the method in relation to the prior work as available in the reviewer comments. However, these responses are short, and the dialogue agent evades the answer by providing a partial/unsatisfactory answer; 2) **Inappropriate answers**: 10% of the errors are incurred because of the lack of knowledge in using comparative adjectives. For instance, responses to questions regarding the confidence of a particular reviewer are often answered with the confidence of all the reviewers. This question is particularly asked as a follow-up to the severity of the weakness stated by the reviewers. However, the dialogue agent answers with the confidence of other reviewers using adjectives such as ‘low’ or ‘high’. These responses are sometimes confusing since confidence levels, such as 4 out of 5, are not particularly low but are conceived as low because of the higher confidence of other reviewers. 3) **Repetition**: 10% of responses have repeated facts as other answers and lack diversity. 4) **Lack of Neutrality**: In 5% of the cases, the dialogue agent agrees with the meta-reviewer’s evaluation of the paper, which is not intended since the dialogue agent should state its neutral stance. However, we did not observe any cases of disagreement with the meta-reviewer, which justifies the refinement metrics used in the dialogue generation.

As a part of our research design, we induce faithfulness and technical specificity in the LLM generations. Thus, errors related to irrelevant answers occur due to increased faithfulness in the generations and the necessity to stick to the source content. The other errors are mainly related to limitations in LLMs in using comparative adjectives and the inherent pre-training-based alignment to agree with humans. However, none of the errors contradict our research design, and all the responses are still technical and faithful. This justifies the use of the dataset to train dialogue agents with higher faithfulness and technicality.

### A.10 Human Evaluation Instruction for evaluating full dialogues

Please evaluate the dialogue in Table 19 with respect to the reviews in Sec §A.24.1 on a Likert Scale of 1 – 4 over the following dimensions:

**Co-operativeness** means that the response of the dialogue agent is coherent with the previous turn and does not try to mislead the interlocutor or act unhelpfully.

**Coherence** how coherent and consistent is the overall dialogue.

**Engagingness** implies the dialogue agent engages the interlocutor by prompting further replies and helps move the conversation forward.

**Plausibility** evaluate the groundedness of the dialogue agent’s response. The measure implies how reasonable the information is about the reviews and the paper as provided by the assistant.

### A.11 Human Evaluation Instruction for evaluating dialogue responses

Please evaluate the response,  $u_t + 1$  in connection to the Knowledge source,  $\mathcal{K}$  and dialogue history  $[u_1, u_2, \dots, u_t]$  on a Likert scale of 1 – 5 over the following dimensions:

**Fluency** How fluent is the language usage in the response?

**Relevance** How appropriate is the response to the dialogue history?

**Faithfulness** Is the response grounded/faithful to the knowledge source,  $\mathcal{K}$ .

**Objectivity** Is the answer factually correct rather than providing personal opinions and feelings?

**Helpfulness** How useful is this response if you had to decide the acceptance or rejection of this paper based on this response?

#### Knowledge Source, $\mathcal{K}$

*Review 1*: This paper proposes a new type of generative models with a new inference method of latent variables. ..., *Review 2*: The paper proposes GONs which seek to build a generative model with an “implicit” encoder that comes essentially for free with the use of a few re-parameterization tricks ..., *Review 3*: This paper introduces a “new” inference method for autoencoder-type models, where the encoder is taken as a gradient of the decoder with respect to a zero-initialized latent variable ...

#### Dialogue History, $[u_1, u_2, \dots, u_t]$

*Meta Reviewer:* Hello, can you summarize the main contributions of this paper "Gradient Origin Networks"?

*Dialogue Agent:* Yes, of course. This paper proposes a new type of generative models called Gradient Origin Networks which use a new inference method of latent variables. Specifically, the gradient of latent variables with respect to the zero vector is taken as the inferred latent variables.

*Meta Reviewer:* I see. Can you tell me more about the pros and cons mentioned in the reviews?]

### **Response, $u_t + 1$**

The reviews describe that the proposed method is easy and straightforward to implement. However, the model assumption that the one step gradient from zero vector equals to latent vector is quite limited and greatly constrains the model expressiveness. Also, the formulation needs to be carefully checked. The experimental results are not sufficient to demonstrate the efficacy. Need more quantitative analysis and experiments on more challenging datasets. The claim that it saves parameters compared to VAE is also confusing.

### **A.12 Annotator Details**

Out of the 15 expert annotators in our study, 7 annotators are 2nd-year Ph.D students while 2 annotators are 1st-year Ph.D students, and the other 6 annotators are 3rd-year Ph.D students. 7 annotators are from India, 6 annotators are from Germany, 1 annotator is from Belgium and 1 annotator is from Italy. All of them are fluent English speakers who have experience in NLP peer-reviewing. The non-expert annotators are from the fields of Mathematics, Railway Engineering, Architecture, Civil Engineering, and Medieval History. Five of the annotators are from Germany, one from Hong Kong, one from India, two from France and the other from Italy.

### **A.13 Correlation between human and automated metrics**

Since human evaluation metrics are widely used and easily interpretable (Macina et al., 2023; Rashkin et al., 2021), we analyze their correlation with automated metrics. We compute correlations between human and automated evaluations (Tables 1 and 2), with results shown in Fig 3. We find that automated metrics closely aligned with human metrics exhibit strong correlations (e.g., K-Prec and Q2: 0.95), while those measuring different as-

pects, such as Specificity and Plausibility, show lower correlation (0.04). Lower correlations within the same category further highlight the diversity of captured aspects.

### **A.14 Disagreement analysis for Human Evaluation in response generation**

We analyzed the instances where the scores varied by 2 or more points (e.g., one annotator scores a response 1 for fluency, while another scores it 4). Out of the 240 responses (60 from each model), the most significant disagreements occurred when rating ChatGPT's responses (12 out of 60). Of these 12 responses, 7 are related to areas of improvement, 4 to the literature review, and 1 to a weakness. All of these responses are notably wordy, averaging 100 words, and often diverge off-topic, a pattern also observed in previous studies (Lozić and Štular, 2023). While these responses are fluent, the variation in objectivity scores between annotators suggests differing perspectives on the content. In contrast, we observed much higher agreement on responses from other models, excluding ChatGPT, with scores for relevance (0.75), faithfulness (0.78), objectivity (0.78), fluency (0.72), and helpfulness (0.70). For ChatGPT, the agreement scores were lower: fluency (0.52), relevance (0.65), faithfulness (0.68), objectivity (0.68), and helpfulness (0.68). Since these measures are used for evaluation rather than to generate gold standards, calibration is performed using averaging, as done in several prior works (Dziri et al., 2022; Rashkin et al., 2021).

### **A.15 Interface for Conversation and Dialogue Agent Evaluation**

We developed the interface using Streamlit to interact with the dialogue agent as shown in Fig 10.<sup>23</sup> In Table 6, we observe higher scores for content relevance and Coverage while using a dialogue agent for both expert and non-expert meta-reviewers. However, decision correctness was lower while using a dialogue agent for the experts, though they perfectly aligned with the non-experts.

### **A.16 Dialogue Process Modelling**

To assess whether ReMuSE captures the decision-making structure of meta-review dialogues, we perform a phase-level evaluation grounded in dialogue-system theory. Each dialogue turn is annotated as one of the three phases, following annotation

<sup>23</sup><https://streamlit.io/>

schemes from task-oriented and debate systems: (i) **Evidence Gathering** – retrieving and summarizing reviewer claims; (ii) **Conflict Resolution** – identifying and weighing contradictory opinions; (iii) **Consensus Recommendation** – synthesizing points of agreement, summarizing trade-offs, and formulating a collective recommendation. We additionally measure **Full 3-Phase Completion**, defined as the proportion of dialogues that include all three phases in the correct sequential order, capturing the model’s ability to simulate a complete meta-review reasoning process.

**Setup.** We train a lightweight RoBERTa (Liu et al., 2019) classifier on 300 randomly sampled manually labeled dialogue turns to identify the three phases, reserving 20% of the data as a hold-out set for evaluation. The classifier achieves a macro-F1 score of **0.83** on this held-out set. For each generated dialogue, we measure (i) the proportion of turns per phase and (ii) phase completeness, defined as the presence of all three phases in sequence.

**Results.** We show in Table 13 that zero-shot dialogues rarely complete all phases (18%), likely because the model lacks structured guidance for meta-review reasoning. Self-Consistency and Self-Refine yield moderate improvements, suggesting that iterative sampling and self-correction provide some structure but remain limited. **We show that ReMuSE achieves the highest coverage**, with 85% evidence gathering, 77% conflict resolution, 81% consensus, and 71% full phase completion, indicating that its training explicitly encourages the model to follow the multi-phase decision-making process and produce more structured, realistic meta-review dialogues.

We show the results for product buying and debates in Tables 14 and 15 respectively. In product buying, dialogues emphasize Evidence Gathering (90%) (researching options) with relatively lower Conflict Resolution (60%), reflecting the focus on exploring alternatives before making a decision. In contrast, debates involve higher Conflict Resolution (counter-arguments) (85%) and moderate Evidence Gathering (78%), consistent with argumentation and synthesis processes. In both domains, ReMuSE consistently achieves the highest phase completion across all phases, demonstrating that modeling structured reasoning improves dialogue quality and generalizes effectively to tasks with different emphasis on information gathering, conflict handling, and consensus-building.

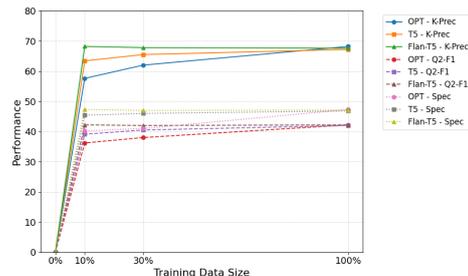


Figure 14: Performance of ReMuSE across different data scales.

### A.17 Effect of data scaling

We show the performance of ReMuSE across different scales : 0,10,30 and 100 % in Fig 14. We find that even with smaller training data subsets (10–30%), all the models achieve substantial gains in K-Precision, Q2-F1, and Specificity, with Flan-T5 already performing competitively at 10% of the data. Increasing the training size to 100% provides only modest additional improvements, suggesting that the algorithmic design and reward-guided fine-tuning, rather than sheer data scale, drive most of the performance gains. This highlights that ReMuSE efficiently leverages limited data to learn knowledge-grounded dialogue behavior.

### A.18 Dialogue Differences across paper types

We show in Figure 15 that dialogue length varies across paper types, with borderline papers exhibiting the longest dialogues (12.2 turns), reflecting the greater deliberation required for conflicting or ambiguous reviews. Accepted papers have moderate lengths (8.3 turns), while rejected papers are shortest (6.1 turns). Figure 16 shows decision correctness, which is highest for accepted papers (90%) and slightly lower for rejected (85%), with borderline papers lowest (78%), indicating that the agent performs best when reviews converge and faces greater challenges in ambiguous cases. Finally, Figure 17 presents phase coverage for evidence gathering, conflict resolution, and consensus formation. Borderline papers have the highest coverage across all phases, especially conflict resolution (77%), emphasizing structured deliberation, whereas accepted and rejected papers show slightly lower coverage. These results highlight that the agent adapts its behavior according to the complexity of the paper, supporting longer and more structured dialogues where necessary.

Model	Evidence Gathering (%)	Conflict Resolution (%)	Consensus Phase (%)	Full 3-Phase Completion (%)
Zero-shot	43 ± 4	28 ± 5	31 ± 6	18 ± 3
Self Con. + Ref.	52 ± 4	36 ± 5	42 ± 5	25 ± 3
Self-Ref.	60 ± 5	44 ± 4	51 ± 5	32 ± 4
<b>ReMuSE (ours)</b>	<b>85 ± 5</b>	<b>77 ± 4</b>	<b>81 ± 5</b>	<b>71 ± 4</b>

Table 13: Modeling of decision-making phases in generated meta-review dialogues. Each column shows the percentage of dialogue turns exhibiting the respective phase. Self-Consistency and Self-Refine provide intermediate improvements over zero-shot, while ReMuSE achieves the highest phase completeness.

Model	Evidence Gathering (%)	Conflict Resolution (%)	Consensus Phase (%)	Full 3-Phase Completion (%)
Zero-shot	55 ± 4	20 ± 5	30 ± 4	18 ± 3
Self Con. + Ref.	65 ± 4	28 ± 5	38 ± 4	25 ± 3
Self-Refine	72 ± 5	35 ± 4	45 ± 5	32 ± 4
<b>ReMuSE (ours)</b>	<b>90 ± 5</b>	<b>60 ± 4</b>	<b>70 ± 5</b>	<b>65 ± 4</b>

Table 14: Phase completion in product buying dialogues. Product buying emphasizes Evidence Gathering (researching options) with fewer conflicts, and ReMuSE achieves the highest completion across all phases.

### A.19 Statistical Measures

We find that paired t-tests and Wilcoxon signed-rank tests, with Cohen’s d, show robust improvements (Cleophas and Zwinderman, 2010). For 15 PhD students, content relevance using the dialogue agent increased by 0.42 ( $t(14)=10.84$ ,  $p \leq .001$ ,  $d=3.43$ ) and coverage by 0.52 ( $t(14)=10.56$ ,  $p \leq .001$ ,  $d=3.34$ ), with task completion time decreasing by 15 minutes ( $t(14)=-19.45$ ,  $p \leq .001$ ,  $d=-6.15$ ). In comparison, Auto Meta-Rev yielded smaller gains (Content Relevance +0.13, Coverage +0.19). For 10 non-PhD participants, content relevance rose by 0.54 ( $t(9)=-13.50$ ,  $p \leq .001$ ,  $d=6.04$ ), coverage by 0.54 ( $t(9)=-10.59$ ,  $p \leq .001$ ,  $d=4.74$ ), and time decreased by 45 minutes ( $t(9)=8.51$ ,  $p \leq .001$ ,  $d=-3.81$ ), while Auto Meta-Rev showed more modest improvements (Content Relevance +0.22, Coverage +0.25). We do observe that the very large effect sizes and **highly significant p-values** indicate these improvements are likely to generalize to a larger population, and that the dialogue agent consistently outperforms both zero-shot and automated meta-review baselines (see Table 6).

### A.20 Performance of other open-source LLMs

We also evaluate additional open-source LLMs, including Qwen 3-8B Instruct (Yang et al., 2025) and Gemma 3-12B-IT (Kamath et al., 2025), with results reported in Table 17 in comparison to baselines discussed in §3.2. Notably, Mixtral’s performance with ReMuSE still surpasses the other models, demonstrating the effectiveness of the dialogues generated by our approach.

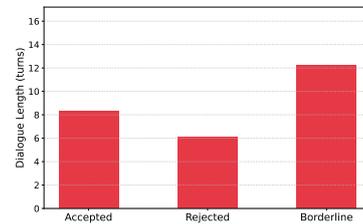


Figure 15: Distribution of dialogue length across paper types.

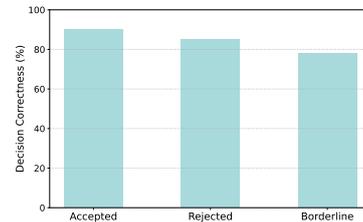


Figure 16: Distribution of decision correctness across paper types.

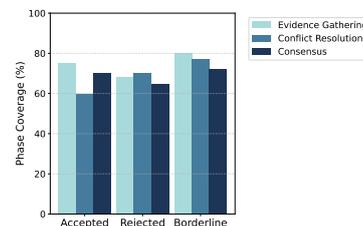


Figure 17: Distribution of phase coverage of dialogues across paper types.

### A.21 Operational guidelines for deploying the dialogue agent

The following operational guidelines outline measures to prevent misuse of dialogue agents in systems that implement our approach presented in Table 16.

Model	Evidence Gathering (%)	Conflict Resolution (%)	Consensus Phase (%)	Full 3-Phase Completion (%)
Zero-shot	35 ± 4	50 ± 5	28 ± 5	15 ± 3
Self Con. + Ref.	45 ± 4	60 ± 5	35 ± 5	22 ± 3
Self-Refine	50 ± 5	68 ± 4	42 ± 5	30 ± 4
<b>ReMuSE (ours)</b>	<b>78 ± 5</b>	<b>85 ± 4</b>	<b>72 ± 5</b>	<b>68 ± 4</b>

Table 15: Phase completion in debate dialogues. Debates involve higher Conflict Resolution (counter-arguments) with moderate Evidence Gathering and Consensus. ReMuSE achieves the highest completion across all phases.

Misuse Prevention Mechanism (Design/UI)	Failure Mode Mitigated	Description
Citation Traceability	“Ghostwriting” / Unverified Claims	UI must display source snippets (reviewer ID, line/context) for all agent-generated claims, enabling instant verification.
Human Final Decision	Automation Bias / Replacement of Human Judgment	The final decision (e.g., “recommend reject”) must be a human-entered field, not an agent-generated output.
Interleaved Presentation	Automation Bias / Reduced Critical Reading	Agent summaries and original source reviews should be presented side-by-side or be easily accessible to encourage comparison.
Balanced Summary View	Anchoring Bias	Agent should present arguments for and against the final recommendation with equal prominence (e.g., randomized order).
Decision Logging	Lack of Accountability / Misuse	Submission process logs whether the agent was used and which agent-generated arguments were directly adopted, creating an auditable trail.
Strict Grounding	Confabulation / Hallucination	Enforce a hard constraint on generation, penalizing statements not mapped to input text; include a UI warning for any unsourced claims.

Table 16: Design/UI mechanisms to mitigate misuse of agent-generated review feedback.

Method	K-Prec	Q2-F1	Spec
<b>Qwen 3</b>			
Zero-Shot	0.36	0.19	0.42
Self-Ref.	0.45	0.20	0.61
Self-Cons. + Ref.	0.74	0.22	0.65
ReMuSE	0.75	0.30	0.72
<b>Gemma 3</b>			
Zero-Shot	0.32	0.17	0.41
Self-Ref.	0.47	0.20	0.62
Self-Cons. + Ref.	0.75	0.21	0.63
ReMuSE	0.75	0.30	0.71
<b>Mixtral</b>			
Zero-Shot.	0.38	0.17	0.42
Self-Ref.	0.50	0.20	0.51
Self-Cons. + Ref.	0.56	0.24	0.59
ReMuSE	<b>0.78</b>	<b>0.32</b>	<b>0.72</b>

Table 17: Performance of different LLMs across various methods. Metrics reported are K-Precision, Q2-F1, and Specificity.

## A.22 End-to-end Workflow

We present the end-to-end workflow in our paper, summarizing all stages in Fig. 18. First, we generate synthetic dialogues for meta-reviewing using our proposed method, **ReMuSE**. Next, we fine-tune dialogue agents to adapt smaller models for the task of meta-reviewing assistance. Finally, we deploy the dialogue agent in real-world settings, demonstrating that this deployment reduces meta-reviewing time by 50%.

## A.23 Dialogue excerpts showing various stages of dialogue agent assistance

In the following table, we show the various stages of dialogue agent assistance.

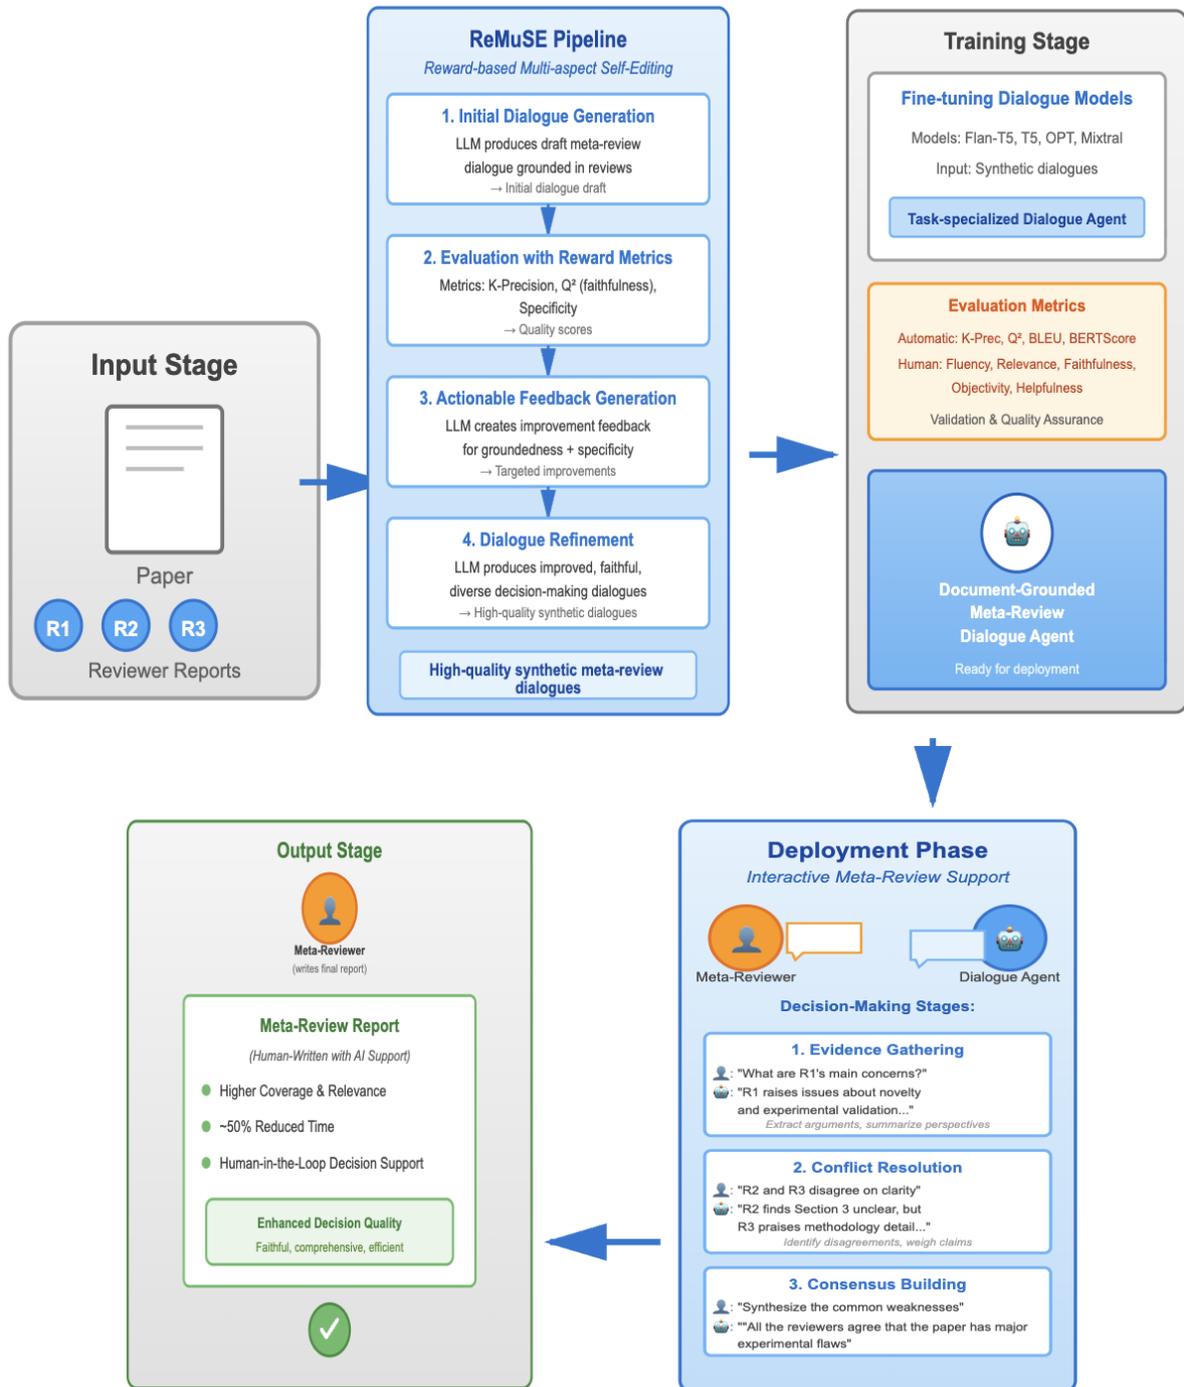


Figure 18: End-to-end diagram outlining the various stages in this work

<b>Role</b>	<b>Dialogue Excerpt</b>	<b>Focus</b>
Meta-Reviewer (MR)	What are the main opposing viewpoints regarding the paper's novelty and empirical support?	Evidence Collection
Dialogue Agent (DA)	Reviewer 1 (R1) rates the paper as Accept (Score: 5) due to "its novel and efficient approach to real-time data streaming that outperforms all listed baselines" [R1, §3]. In sharp contrast, Reviewer 3 (R3) recommends Reject (Score: 2), arguing that the "empirical gains are only marginal on standard benchmarks and the proposed method introduces too much complexity" [R3, §2].	Conflict Identification
MR	R3's point about "marginal gains" and complexity is concerning. Does R1 or R2 offer any perspective on why the complexity is justified?	Resolution Inquiry
DA	Reviewer 2 (R2), who rated the paper 4/Weak Accept, directly supports the added complexity. R2 states, "The conceptual elegance of the attention mechanism alone justifies the extra parameters, as it opens up new avenues for interpretability" [R2, §4]. This suggests the complexity is a trade-off for a valuable architectural feature.	Evidence Retrieval / Conflict Mitigation

Table 18: Dialogue excerpts illustrating meta-reviewing process.

## A.24 Example Dialogues from our Dataset

We show a few reviews from our dataset along with the dialogue generated using ReMuSE below.

### A.24.1 Meta-Reviewing

The title of the paper is not revealed to retain anonymization. The type of the paper is “long”.

**Review 1:** Paper is on modeling the prediction of ancestor relation between names of science institutions. This is on the GRID dataset which already has some hierarchical information. The proposed approach is set-based models (with neural encodings) where the overlap between two names is measured by set overlap at the unigram level. In extended experiments additional metadata like address and type of institution are also incorporated into the model (which contribute a lot to the improvements). A set of simple to intermediate baseline along with different thresholds of token overlap has been tested and the proposed model shows strong improvement in the MAP metric. Paper has a decent writing and structure. Problem and the approach has been explained and motivated well with descriptive examples. [...]

**Review 2:** The paper shows how to infer the organisational structure of an institution. That is, it presents a model for predicting the is-ancestor relationships of institutions based on their string names. To this end, it makes use of Set-Transformers to model the token overlap between the institution names. This use is nice but also not highly original. The experimental evaluation is on a single dataset only. While the authors do present some examples, and overall hierarchy or something that provides some more insights into the learned model should be provided in order to show potential issues with transitivity and connected components. The evaluation only considers known pairs. But an organisational structure should also be consistent. That is, the interesting motivation provided in the intro is not met in the experimental evaluation. Furthermore, the experimental protocol is unclear [...]

---

<b>Meta-reviewing Dialogue</b>	
<i>Meta-Reviewer</i>	Hello, I'm reviewing a paper on predicting institution hierarchies with set-based models. Can you tell me a little bit about the paper?
<i>Dialogue Agent</i>	Sure! The paper presents a new approach to predicting the hierarchical structure of institutions using set-based models with neural encodings. The approach is based on the GRID dataset, which already has some hierarchical information.
<i>Meta-Reviewer</i>	That sounds interesting. Can you tell me more about the dataset and the experimental setup?
<i>Dialogue Agent</i>	Sure! The dataset is the GRID dataset, which is a global research identifier database. The authors have constructed a new dataset by pairing institution names with their corresponding hierarchy information. They have also compared their approach against a few baseline approaches and presented some examples.
<i>Meta-Reviewer</i>	Okay, that sounds good. But I have a few doubts regarding the experimental setup. For example, I would have expected a comparison with well-known baseline models from the knowledge base completion literature, such as TransE. And I'm not sure if the manually extracted test set is complete. Can you tell me more about this?
<i>Dialogue Agent</i>	Yes, the authors have mentioned some of these concerns in their paper. They have developed a new dataset, and the task is well-defined. However, they have not compared their approach with well-known baseline models from the knowledge base completion literature, such as TransE. And the manually extracted test set may not be complete, which may affect the evaluation.
<i>Meta-Reviewer</i>	I see. That's a good point. What about the token-based models? Are they a promising choice for the task?
<i>Dialogue Agent</i>	The authors have presented some examples in Table 3 to show the effectiveness of their approach. However, the task is presented as a knowledge-base completion task, and the token-based models may not be the best choice for this task.
<i>Meta-Reviewer</i>	Okay, I understand. What about the related work on set-based models? Is there any extension of the related work in the paper?
<i>Dialogue Agent</i>	The authors have extended the related work on set-based models in their paper. They have presented some new approaches and compared them with the baseline models. However, the related work on set-based models should be extended further.
<i>Meta-Reviewer</i>	I see. Thank you for the information. Based on your provided information, I would rate this paper as a 6 out of 10, with a confidence level of 4. The paper has some interesting ideas and the authors have presented some examples to show the effectiveness of their approach. However, there are some concerns regarding the experimental setup and the token-based models. The related work on set-based models should be extended further.
<i>Dialogue Agent</i>	Thank you for your feedback.

---

Table 19: Generated Meta-Reviewing Dialogue from our dataset based on the provided reviews.

### A.24.2 Debates

The topic of the paper is: 'Global Warming is Not a Crisis'. The arguments for for and against sides of the debate are provided below:

**Against side:** "Global warming is not a crisis. Please keep in mind what the proposition is. It is not a debate over whether the earth has been warming over the past century. Uh, the earth is always warming or cooling, at least a few tenths of a degree. And we're talking about, so far, something on the order of six tenths of a degree centigrade. We're not even arguing about whether greenhouse gas emissions are contributing at some level to warming. And they most certainly should or I would suggest it would be very little. Indeed, as far as I can tell, even our opponents do not claim that global warming is a crisis at present. Rather, we are primarily addressing the future. Now, much of the current alarm, I would suggest, is based on ignorance of what is normal for weather and climate. [...]"

**For side:** "The motion before us, global warming is not a crisis, means we ought to know what crisis means. The word does not mean catastrophe or alarmism. It means a crucial or decisive moment, a turning point, a state of affairs in which a decisive change for better or worse is imminent. We are talking about the future here. The entire world now really does have a critical choice to make. It is whether to continue on the present path of adding more and more carbon dioxide and other greenhouse gases to the atmosphere or whether to find another path. We're speaking of the future. And science tells us that the path we choose will largely determine what kind of earth our children and grandchildren will inherit. Our task tonight is to persuade you that global warming is indeed a crisis in exactly that precise sense so you should vote against the motion. The science community today has impeccable settled science, despite what you have just heard, that demonstrates the reality of global warming and its primary origin in human activities. [...]"

---

**Debate Decision-Making Dialogue**

---

<i>Decision Maker</i>	Hello, I am here to learn more about the debate topic, "Global warming is not a crisis." Can you provide me with information about both sides of the argument?
<i>Dialogue Agent</i>	Of course, for the "For" side, they argue that global warming is not a crisis at present and much of the current alarm is based on ignorance of what is normal for weather and climate. They suggest that extreme weather events always occur without evidence of systematic increases. Additionally, the sea level has been increasing since the end of the last Ice Age glaciation. The impact of warming on agriculture and disease seems dubious at best. On the other side, the "Against" arguments state that global warming is indeed a crisis in the sense that it is a turning point, a state of affairs in which a decisive change for better or worse is imminent. They argue that science has proven the reality of global warming and that mankind's fingerprints have now clearly emerged above the noise of natural variability. They also mention the increasing frequency of heat waves, higher sea levels, and vanishing glaciers as evidence of the severity of the situation.
<i>Decision Maker</i>	Thank you for providing that information. Can you tell me about any projections or forecasts related to global warming?
<i>Dialogue Agent</i>	Yes, the "Against" side states that they project further warming of half a degree Fahrenheit for the next twenty-five years. Beyond that, it largely depends on how much more CO <sub>2</sub> and other greenhouse gases humanity dumps into the atmosphere. They also warn that continuing to fuel the world using present technology will bring dangerous and possibly surprising climate changes by the end of this century, if not sooner.
<i>Decision Maker</i>	What is the scientific community's stance on global warming?
<i>Dialogue Agent</i>	The "Against" side argues that today's science community has impeccably settled science that demonstrates the reality of global warming and its primary origin in human activities. They also mention the Intergovernmental Panel on Climate Change's report and the extensive process of experts who represent mainstream science. On the other hand, the "For" side acknowledges that even their opponents do not claim that global warming is a crisis at present and suggest that the impact of CO <sub>2</sub> per unit of CO <sub>2</sub> actually goes down with increasing CO <sub>2</sub> . They also argue that it is crucial to distinguish between the claim that models can display past behavior from the actual situation, which is that models can be adjusted to display past behavior once that behavior is known.
<i>Decision Maker</i>	Thank you for the information. Based on both sides of the argument, what do you suggest is the conclusion?
<i>Dialogue Agent</i>	As a dialogue agent, I cannot express opinions or make decisions. However, I provided the arguments from both sides that may help you reach your own conclusion.

---

Table 20: Debate Dialogue on the topic 'Global warming is not a crisis' generated from the provided arguments.

### A.24.3 Product Reviews

The product in discussion is "Flexion KS-902 Kinetic Series Wireless Bluetooth Noise Cancelling Headphones with Microphone, Black". The reviews are below:

**Reviews:** "The battery life is sufficient where I only need to charge once every 2 days These headsets are equipped with the new Bluetooth 9 edition too advance for current music, so don't even bother to enable because they don't work. \n

These gave very good sound. Good sound quality, fits pretty well in your ear And it is also bluetooth, so I can connect it to my phone easily!\n

The user interface is really bad, when I try to turn it off, it usually calls someone on my phone, so I have had several butt dials thanks to these earphones.\n

There well rounded in all areas something that others don't have. And they totally out perform the 2 hour battery life of the Flexion. The case is badly designed and too small.\n

These work really great when you want to listen to something without bothering other people. sleep.- Bass is amazing and really good- Loud and really clear For the price I paid they work excellent!\n

my background noise was amplified. He told me that he loves the headphones, but it was the amazing C.S that will have him as a returning customer. I recommend it to everyone! They actually turn on with a quick click, not a 3 second hold as in the instructions. The sound is just fantastic, the best sound I ever hear from an ear bud type speaker. Then my replacement pair broke, the cord just separated from the ear bud during normal use (see picture attached). \n

Sound quality is good and has buttons on the ear pieces for volume control, answering phone calls, and changing songs! The pause button and the advance track button doesn't work and the previous track button is awkward and works only some of the time. As a commuter using public transportation, I was looking for something that would eliminate wires going to to my device and these do a great job of that. build quality is great compared to the price they do look expensive\* I love that they were really easy to sync, which is a huge bonus for me! The headphones connected to my iPhone with no problems. \n

Sound quality was great but even after a full night of charging the batteries died after 2 hours. they give you a low battery warning a few seconds before they shut off which is kinda pointless. \n [...]

<b>Product Buying Dialogue</b>	
<i>Buyer</i>	Hi, I'm interested in purchasing the Flexion KS-902 Kinetic Series Wireless Bluetooth Noise Cancelling Headphones with Microphone, Black. Can you provide me with some information about the product?
<i>Dialogue Agent</i>	Sure, what would you like to know?
<i>Buyer</i>	Can you tell me about the battery life of the headphones?
<i>Dialogue Agent</i>	According to the reviews, the battery life is sufficient where you only need to charge once every 2 days. However, there have been a few reviewers who experienced a shorter battery life.
<i>Buyer</i>	What about the sound quality?
<i>Dialogue Agent</i>	The reviews state that the sound quality is good, and they worked really great when you want to listen to something without bothering other people. However, there were some reviews that reported the sound quality wasn't that great, and some experienced full of static if you get more than 7 feet from your phone.
<i>Buyer</i>	Is it easy to connect to my phone?
<i>Dialogue Agent</i>	Yes, the headphones are equipped with Bluetooth and are easy to sync with your iPhone or Android device.
<i>Buyer</i>	What about the user interface? Is it easy to use?
<i>Dialogue Agent</i>	According to some reviews, the user interface is really bad. They experienced several butt dials when trying to turn it off. However, others found the headphones to be easy to connect and use.
<i>Buyer</i>	With all that information, do you recommend these headphones?
<i>Dialogue Agent</i>	I'm sorry, I cannot provide recommendations or opinions. Please use the reviews provided to make an informed decision.
<i>Buyer</i>	Okay, thanks for the information. I think I'll go ahead and purchase them!

Table 21: Dialogue generated for the product 'Flexion KS-902 Kinetic Series Wireless Bluetooth Noise Cancelling Headphones with Microphone, Black' with the reviews provided before.